

Issue 2, 1994

Personal Software

Software

- 1** What's New in PC DOS 6.3
- 13** PC DOS 6.x Disk Compression & Defragmentation
- 25** OS/2 Workplace Software Strategy

Personal Software Magazine



IBM Personal Software Magazine is published by IBM Personal Software Products User Group Relations, Austin, Texas, U.S.A.

User Group Relations Program Manager

Gene Barlow

Editorial Services

MDE Associates

Illustrations

Jamison American Originals

Typesetting

Dove Oaks Publishing

This publication is distributed only in bulk to PC and OS/2 user groups. Individual copies and subscriptions are not available.

Correspondence should be addressed to Gene Barlow, IBM Corporation, P.O. Box 201449, Austin TX 78720-1449. If you're electronically connected, you can send an Internet note to IBMPCUG@VNET.IBM.COM. IBM employees can send to IBMPCUG@AUSVM1.

IBM believes the statements contained herein are accurate as of the date of publication of this document. However, IBM hereby disclaims all warranties as to materials and workmanship, either expressed or implied, including without limitation any implied warranty of merchantability or fitness for a particular purpose. In no event will IBM be liable to you for any damages, including any lost profits, lost savings or other incidental or conse-

quential damage arising out of the use or inability to use any information provided through this service even if IBM has been advised of the possibility of such damages, or for any claim by any other party.

Some states do not allow the limitation or exclusion of liability for incidental or consequential damages so the above limitation or exclusion may not apply to you.

This publication could contain technical inaccuracies or typographical errors. Also, illustrations contained herein may show prototype equipment. Your system configuration may differ slightly.

IBM has tested the programs contained in this publication. However, IBM does not guarantee that the programs contain no errors.

This information is not intended to be a statement of direction or an assertion of future action. IBM expressly reserves the right to change or withdraw current products that may or may not have the same characteristics or codes listed in this publication. Should IBM modify its products in a way that may affect the information contained in this publication, IBM assumes no obligation whatever to inform any user of the modifications.

Some of the information in this magazine concerns future products, or future releases of products currently commercially available. The description and discussion of IBM's future products, performance, functions, and availability are based upon IBM's current intent and are subject to change.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not imply giving license to these patents.

It is possible that this material may contain reference to, or information about, IBM products (machines and programs), programming or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such products, programming or services in your country.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation whatever.


All specifications are subject to change without notice.

© Copyright 1994 by International Business Machines Corporation.

What's New in PC DOS 6.3

Robert Rapuano
IBM Corporation
Boca Raton, Florida
and

Nestor Miranda
Electronic Data Systems
Boca Raton, Florida



*PC DOS 6.3 is the latest release of IBM's PC DOS for IBM and IBM-compatible personal computers. Its features and functions are the cumulative result of many years of user feedback plus innovative design. This article presents an overview of the features of PC DOS 6.3. Such an overview necessarily includes discussions of features carried over (and enhanced) from PC DOS 6.1 and previous versions of DOS. The features that are new in PC DOS 6.3 are appropriately identified with the word **NEW!***

In a Nutshell

Here is a quick summary of what is new in PC DOS 6.3:

- **Smart Upgrade** from MS-DOS** 6.0/6.2 or PC DOS 6.1 to PC DOS 6.3. Install has been modified to understand the multi-configuration support delivered in these versions of DOS, making upgrading easier and safer.
- **RAMBoost** now searches the CONFIG.SYS file to determine whether DOS 6 multi-configuration has been used. If found, RAMBoost** has the ability to optimize and track each environment separately, giving the user dynamically optimized memory utilization in multi-configuration environments. RAMBoost now also has

built-in DOS-UMB support, which (when activated) advises DOS that it is running in DOS-UMB mode. This means that direct XMS-UMB memory requests and LOADHIGH commands will be honored from the command line.

- **Backup** now supports more tape drives.
- **Compression** has been fully integrated.
- **SMARTDRV** has been improved, and will now cache CD-ROMs.
- **DEFRAG** now uses extended memory to efficiently defragment even large hard disks easily.
- **QCONFIG** diagnostic tool has been updated to recognize a host of new OEM adapters. It also comes with an enhanced reporting scheme.
- **MSCDEX 2.23** provides access to CD-ROM drives.
- **E Editor** has been enhanced with a new command (ALL) that creates a separate file of all search instances found in a document.
- **AntiVirus** has been enhanced to detect and eliminate more than 2,000 viruses.
- **Thousands-digit separator** now appears when using DIR, CHKDSK,

FORMAT, and MEM commands. This makes numbers easier to read. (The display of the thousands-digit separator can be turned off.)

- **No-Swap DISKCOPY** copies the entire contents of one floppy diskette to another by using the system's hard disk as a temporary storage area. The laborious task of swapping diskettes multiple times is no longer necessary.
- **Interactive Boot** lets you bypass or step through your CONFIG.SYS and AUTOEXEC.BAT files. It lets you start with a clean system without having to rename your boot-up files. The step-through option makes it easier to identify where you may be having a problem booting up your system.
- **Interactive Batch Processing** brings the step-through feature to any batch file.
- **File Overwrite** prompts you before overwriting files that have identical file names when you use the COPY, XCOPY, or MOVE command. Overwrite protection can be easily turned off.
- Keyboard support has been added for Macedonia and Serbia (Cyrillic).

Figure 1 compares major features of PC DOS 6.1, PC DOS 6.3, MS-DOS 6.2, and MS-DOS 6.21.

	PC DOS 6.1	PC DOS 6.3	MS-DOS 6.2	MS-DOS 6.21
Memory Management				
Dynamic memory optimizer	Yes	Yes	No	No
Automatically loads programs and drivers in UMBs	Yes	Yes	Yes	Yes
UMB device drivers for 8088/80286	Yes	Yes	No	No
Tracks and optimizes multiple configurations	No	Yes	No	No
Advanced upper-memory usage editor	Yes	Yes	No	No
Optimizer supports memory managers other than EMM386	Yes	Yes	No	No
Disk Compression				
Automated compression uninstall routine	Yes	Yes	Yes	No
Universal Data Exchange (transportable floppies)	Yes	Yes	No	No
Auto-mount compressed floppies	Yes	Yes	Yes	No
Active compression preview	Yes	Yes	No	No
Bypass compression on startup	Yes	Yes	Yes	No
DoubleGuard	No	No	Yes	No
Password protection	Yes	Yes	No	No
Compression-aware utilities (FORMAT, UNFORMAT, DISKCOMP, DISKCOPY, SYS)	Yes	Yes	No	No
DoubleSpace-compatible	Yes	Yes	Yes	No
Backup				
Tape support	Yes	Yes	No	No
Restores files from DOS command line	Yes	Yes	No	No
Uses data compression	Yes	Yes	Yes	Yes
File viewers	Yes	Yes	No	No
High-speed floppy disk support	Yes	Yes	No	No
Unattended (scheduled) backup	Yes	Yes	No	No
Three levels of user interface	Yes	Yes	No	No
DOS and Windows interface	Yes	Yes	Yes	Yes
AntiVirus				
Fuzzy logic detects new variants	Yes	Yes	No	No
Avoids false alerts by other anti-virus products	Yes	Yes	No	No
Scans Boot Manager boot records for viruses	Yes	Yes	No	No
Integrity checking	Yes	Yes	Yes	Yes
Signature scanning	Yes	Yes	Yes	Yes
Configurable alert messages	Yes	Yes	No	No
Option to prevent user from disinfecting	Yes	Yes	No	No
Shield and scanning options	Yes	Yes	Yes	Yes
Viruses detected	1,500	2,000	800	800
RAM requirements	1-6KB	1-6KB	7-44KB	7-44KB

Figure 1. Comparison of Features in PC DOS 6.x and MS-DOS 6.x

	PC DOS 6.1	PC DOS 6.3	MS-DOS 6.2	MS-DOS 6.21
Editor				
Edit multiple files simultaneously	Yes	Yes	No	No
String search/replace	Yes	Yes	Yes	Yes
Reflow paragraph	Yes	Yes	No	No
Mouse-aware	No	No	Yes	Yes
Undo	Yes	Yes	No	No
Split-screen view	Yes	Yes	No	No
Draw lines and boxes	Yes	Yes	No	No
Sort, add, or multiply within a marked area	Yes	Yes	No	No
PC Management Tools				
No-swap DISKCOPY	No	Yes	Yes	Yes
Interactive AUTOEXEC.BAT	No	Yes	Yes	Yes
MOVE, COPY, XCOPY confirmation on overwrite	No	Yes	Yes	Yes
SMARTDRV as a read-only cache by default	Yes	Yes	Yes	Yes
Defragger	Yes	Yes	Yes	Yes
Enhanced defrag	No	Yes	Yes	Yes
ScanDisk	No	No	Yes	Yes
Data Recovery				
Undelete Windows user interface	Yes	Yes	Yes	Yes
Full-screen undelete DOS user interface	Yes	Yes	No	No
Undelete directories (from DOS interface)	Yes	Yes	No	No
Undelete file viewers	Yes	Yes	No	No
Advanced undelete methods	Yes	Yes	No	No
Other Utilities				
Program scheduler	Yes	Yes	No	No
PCMCIA support	Yes	Yes	No	No
Pen support	Yes	Yes	No	No
QBASIC	No	No	Yes	Yes
DOS Shell with task switching	Yes	Yes	No	No
APM Power Management	Yes	Yes	Yes	Yes
Interlink file transfer	Yes	Yes	Yes	Yes
Miscellaneous				
Command reference	Hard copy	Hard copy	On line	On line
Comprehensive printed documentation	Yes	Yes	No	No
MSCDEX (CD-ROM support)	No	Yes	Yes	Yes
ISO font support	Yes	Yes	No	No
Smart upgrade (multi-configuration-aware installation)	No	Yes	No	No
Display of thousands separator	No	Yes	Yes	Yes

Figure 1. (continued) Comparison of Features in PC DOS 6.x and MS-DOS 6.x

Memory Management

EMM386. PC DOS 6.3's memory manager is EMM386 (not to be confused with the optimizer RAMBOOST).

EMM386 has the ability to make use of available Upper Memory Block (UMB) RAM. Its search algorithm finds more free space in High Memory Area (HMA) as well as in unallocated areas of memory.

Another improvement in EMM386 is its ability to allocate expanded and extended memory from shared pools; EMM386 no longer has to choose between EMS and XMS when starting your system.

UMB Device Driver Support. PC DOS 6.3 includes additional UMB device driver support for many users who have 8088 and 286 computers. Many of these computers might have been upgraded to support EGA/VGA displays, or they might have had EMS boards installed to run memory-intensive applications.

UMB Device Driver Support enables users who have video cards or EMS boards to obtain up to an additional 60 KB of free memory. With PC DOS 6.3, the extra memory found in these boards can be utilized.

NEW! Multi-Configuration Option Aware. PC DOS 6.3 searches the CONFIG.SYS file to determine if the DOS multi-configuration option is used. If found, PC DOS 6.3 will optimize and track each environment separately, giving the user dynamically optimized memory utilization in multi-configuration environments. (In previous versions of DOS, neither MEMMAKER nor RAMBOOST would properly optimize systems that were taking advantage of DOS's multi-configuration option.)

Memory Optimizer (RAMBOOST). PC DOS 6.3's optimizer attempts to find the best arrangement for placing device drivers and Terminate and Stay Resident

(TSR) programs into UMB. By getting as much code as possible into upper memory, the optimizer lets your DOS applications (which run only in 640 KB) run faster and more efficiently.

The optimizer in PC DOS 6.3 works the same as most other optimizers. It first scans your configuration files and determines which TSRs and device drivers are being loaded; then, it probes memory and orders the TSRs and device drivers in the predetermined optimal way; finally, it saves the optimized configuration and restarts your computer.

Configuration Watcher. PC DOS 6.3's memory optimizer has an intelligent learning and tracking feature that watches over your systems configuration. If modifications have been made to your configuration (such as adding programs), the configuration watcher will prompt you. This type of optimization is called *install and forget* or *dynamic* memory optimization.

This optimization is different from other memory management programs. For example, in MS-DOS 6.2, MEMMAKER uses a "static" snapshot approach — when MEMMAKER is executed, it optimizes your configuration; however, it will not reoptimize unless told to do so by the user. This approach assumes the user is always aware when changes are made to the system; in truth, however, most users are not aware of these changes, so optimization degrades over time, and the system will not make best use of upper memory.

For example, suppose you installed a sound card and a CD-ROM drive in your system. When you reboot your system, RAMBOOST displays a message saying that changes have been made, and asks if you would like to re-optimize. After you select Yes, the system optimizes itself.

In previous versions of IBM DOS, and in other versions of DOS, you would

have had to update your configuration manually to reflect the most efficient way to load your new drivers. Also, if you were using MS-DOS's MEMMAKER, you would have to remember to rerun the optimizer.

Advanced Upper Memory Editor. PC DOS 6.3 includes an Advanced Upper Memory Usage Editor that gives you an easy, graphical way to directly manipulate HMA contents.

NEW! UMB Support. PC DOS 6.3's memory management comes with built-in DOS-UMB support. When activated, the memory manager advises DOS that it is running in DOS-UMB mode, causing direct XMS-UMB memory requests and LOADHIGH commands to be honored from the command line.

Disk Compression (SuperStor/DS)

Disk Compression. Safe and reliable file compression has been integrated into PC DOS 6.3. PC DOS's compression works with both DOS and Windows** programs. Compression in PC DOS 6.3 is compatible with the DoubleSpace compression program in MS-DOS 6.2, and compressed files in PC DOS 6.3 work the same as they do with DoubleSpace. You do not have to remove the DoubleSpace compression program, or uncompress any DoubleSpace compressed volume files, in order to replace them with PC DOS 6.3's compression (SuperStor/DS). All you need to do is select compression and install it.

An additional feature in PC DOS's compression is its ability to load the compression driver prior to loading the CONFIG.SYS. This feature eliminates the need for multiple CONFIG.SYS files, and makes compression relatively transparent to the end user.

Once installed, PC DOS 6.3's compression effectively increases your hard-disk

capacity by compressing data up to an average of 50 percent. For example, suppose your PC has an 80 MB hard disk with only 10 MB of free space. After you compress the data on your hard disk, its capacity effectively doubles to 160 MB (the actual ratio varies depending on data). In previous versions of IBM DOS, and in other versions of DOS, you would have had to either maintain data files on diskettes, delete programs not used often, or purchase a larger disk drive.

By making PC DOS's compression compatible with MS-DOS compression, upgrading to PC DOS 6.3 is safe and easy.

Universal Data Exchange. This valuable feature makes compressed disks portable within your environment. It enables a diskette containing compressed data, and created on one computer, to be used on another computer, whether or not both computers are running compression software. A disk compressed with PC DOS 6.3 can be read and written by a computer running *any* version of DOS (PC DOS, MS-DOS or DR-DOS).

As an example, suppose that in your office your computer is the only one running compression software. Your accounting file uses 2.4 MB, which does not fit on a single diskette; but, because you are running PC DOS 6.3 with compression, the accounting file fits easily on a single compressed diskette. Now suppose your officemate needs the information in your accounting file. You give her the diskette; she inserts it into her computer and types UDEON. The information on your compressed diskette is now fully accessible – transparently – even though your officemate is running another version of DOS without compression software. Your officemate brings up her accounting package, reads the compressed data, and writes changes back to the disk (in compressed format).

Auto-Mounting Diskettes. The auto-mounting diskettes feature gives the

File Type	Range of Probable Compression Ratios
Executable programs (.EXE)	1.4:1 to 2:1
Word processing documents	2:1 to 4:1
Text files	2:1 to 8:1
Database files	2:1 to 8:1
Spreadsheet files	2:1 to 4:1
Video image files	2:1 to 8:1
CAD/CAM files	3:1 to 8:1
.ZIP or .ARC files	Will not compress further

Figure 2. Compression Ratios for Different Kinds of Files

computer automatic access to compressed diskettes. Without it, the user would have to enter commands to enable access.

Automated Deinstall makes it easy to convert a compressed hard disk back to its original uncompressed form. A diskette, however, cannot be uncompressed by PC DOS 6.3.

Compression Preview. PC DOS 6.3 comes with compression preview. This feature scans your hard disk and then estimates (based on what kind of data is on your hard disk) how much your data will compress. Compression ratios vary, depending on which data type is being compressed.

For example, suppose that when you want to compress your hard disk, you are asked to enter a compression ratio between 2:1 and 8:1. Not certain about which ratio to specify, you enter the default, 2:1. However, most of your files are graphic and database files, and could actually be compressed to 5:1. Here, specifying the default value means that more disk space is being used for compression than is actually needed. To prevent wasting valuable disk space, you can run the compression preview, which recommends a 5:1 ratio. Without this feature, you would have to guess which ratio to enter.

Figure 2 is a table of the most likely compression ratios for different kinds of files.

Backup (Central Point)

Backup. PC DOS 6.3 provides a full-featured DOS and Windows version of Central Point's PC Tools Backup** utility program. Both user interfaces are easy and intuitive, using a tree-view display of your file system for easy point-and-click operations. If you do not want to use the DOS or Windows interface, PC DOS 6.3 lets you run Backup from the command line.

File Viewers. PC DOS 6.3's Backup comes with a set of valuable file viewers that let you view the contents of a file to determine if you want to back it up. Both DOS and Windows file viewers are included.

Scheduler. PC DOS 6.3's Backup program includes an easy-to-use scheduler program. Because backing up data can take valuable time, you can use this feature to specify that the backup should occur during a time when the computer is idle.

User Levels. PC DOS 6.3's Backup program lets you select among three different levels of backup. You can configure Backup for entry-level, experienced, or advanced usage.

Backup Media. PC DOS 6.3's Backup lets the user specify the type of backup (full, partial, or incremental) and a variety of backup media. In addition, PC DOS's Backup utility compresses the data while backing up, thus saving valuable space on the target media. Because many hard disks exceed 200 MB, a full backup solution includes the ability to backup to tape drives. Unlike MS-DOS, PC DOS provides support for many tape drives.

The following media are supported by PC DOS 6.3 Backup:

- Tape
- Optical
- Hard disk
- Diskette
- High-speed diskette

Online Help. PC DOS 6.3's Backup comes with a comprehensive, easy-to-use online help.

AntiVirus (IBM Research)

AntiVirus was developed by IBM Research. With a full-screen DOS and Windows interface, it is an industrial-strength product that is highly competitive with current industry leaders, and is used by IBM and some of its largest accounts.

AntiVirus detects and cleans both known and unknown viruses; it finds over 2,000 viruses, including generic lab, polymorphic, and stealth viruses. (Generic Lab viruses are created with a specific virus creation tool; polymorphic viruses can change themselves; stealth viruses take actions that try to protect themselves from virus detection.)

AntiVirus does not damage the programs that it is disinfecting. It disinfects the program only when it knows exactly what the virus is, and how to disinfect it properly. If it does not recognize the virus, it notifies you of its suspicion.

PC DOS 6.3's AntiVirus has been developed as an "install-and-forget" protection system for your computer. One of its components is active all the time, preventing common viruses from becoming active in your system and spreading. As soon as it is installed, AntiVirus not only protects the DOS sessions in which you run programs, but will also run automatically at prescheduled times to check your entire disks for viruses and allow you to disinfect them.

AntiVirus follows the principle that anti-virus software should be unobtrusive. It runs in the background, even without your knowledge. You do not have to know anything about the AntiVirus program; you need not become a virus expert; and you need not be a computer whiz in order to use AntiVirus.

AntiVirus employs three virus detection schemes:

Integrity checking is often used to try to detect unknown viruses. It tries to find out when things have changed in the system, and if these changes have been done by a virus rather than something legitimate. It will tell you that changes have occurred, but you must determine whether they were caused by a virus. In truth, most users have trouble making this determination; they often don't know when to expect changes, and when not to.

PC DOS 6.3's AntiVirus recognizes how viruses change things. It only tells you when things change in a way that appears to have been caused by a virus. This method drastically cuts down on the number of false alerts that you get, and leads you to have more trust in your virus program when it does issue an alarm.

Signature scanning is a technology for finding known viruses. It looks for a sequence of bytes, or something a little more complicated, that exists in a known virus. If a certain sequence of

bytes occurs in a virus, and that sequence appears in a file or a boot sector on a computer, then there is a virus in that computer.

PC DOS 6.3 has technology that avoids false positives or false alarms. PC DOS 6.3 does an exact verification of all viruses that are common to DOS. However, it tells you about a virus only if its sequence matches byte-for-byte.

Fuzzy logic detection is a technology that IBM developed in signature scanning. Typical signature scanners look for some number of bytes; if the virus detector recognizes such a string, it indicates the presence of a virus. Many viruses are variants of older viruses wherein a few bytes have been changed. This technique is used to circumvent anti-virus program signature strings – modify one byte in the middle of a signature string, and the anti-virus program cannot detect it.

With PC DOS 6.3 AntiVirus, a fuzzy scanning technique is used. This technique is smart enough to determine when these changes have occurred, and is less prone to allow a virus to go by undetected.

TSR Protection. PC DOS 6.3's anti-virus protection comes in the form of a TSR program. This is crucial, because an effective protection scheme requires an active sentry at all times to warn you when a virus is present or when one has been introduced to your system. There are a number of technologies for doing this. PC DOS employs a unique technology that can tell you when you are executing a program that contains any of the common viruses. It will stop the virus from propagating, which in turn will keep the virus from doing damage.

Run Infected Programs. Another valuable feature in PC DOS 6.3's AntiVirus program is its ability to safely run infected programs as though they are not infected. This ability to run critical appli-

cations that have been infected by a virus, without any fear of spreading the infection or doing damage, is indispensable.

Virus Service Plan. PC DOS 6.3's virus update plan is a flexible service plan that includes virus updates. IBM has extensive knowledge about dealing with viruses, and offers support that ranges all the way up to fee support for the customer who has a virus incident and needs assistance managing it. In this service plan, users receive updates, bulletin board access for questions, and access to the Virus Information Manual (a comprehensive manual about virus incident management and descriptions of viruses). If a user happens to get a virus that has not been seen before, the user can get rapid response from IBM AntiVirus Services, which will give the user information for detection and ways to eliminate the virus as quickly as possible.

RAM Requirements. PC DOS 6.3's AntiVirus memory footprint is very small compared to those of other anti-virus programs. It is important for anti-virus software to have as small a memory profile as possible. When loaded into expanded memory, PC DOS 6.3's AntiVirus program takes up as little as 1 KB; when loaded into conventional memory, the anti-virus program takes up to 6 KB. (In contrast, MS-DOS's memory footprint can take as much as 44 KB.) Because its memory footprint is small, PC DOS 6.3 allows users to keep the AntiVirus program running at all times, ensuring maximum protection while not using memory that can be better utilized by other programs.

Network Drives. Different anti-virus software packages scan network drives with one program and local drives with another. PC DOS 6.3's AntiVirus program scans all drives at once, no matter how they are configured. In short: If DOS recognizes the drives, PC DOS

6.3's AntiVirus program will protect them.

Customized Alert Messages. PC DOS 6.3's users can personalize alert messages to say anything they want. The system administrator can customize alert messages for specific phone numbers and directions to follow when a virus is encountered. For example, you could display a message that says "Warning: Contact Virus Specialist Elaine Smith at extension 675."

E Editor (IBM Research)

The E Editor was designed to be a full-function editor, providing all the basic editing functions you need. Its power and flexibility enables you to do many of the editing functions found in high-end word processors.

With PC DOS 6.3's E Editor, you can edit multiple files, view them simultaneously, cut and paste between them, set margins and tab stops, re-flow paragraphs, search and replace, and undo previous actions. The E Editor supports split-screen views, so you can view several files at the same time. It has an autosave feature that helps ensure you don't lose the data you have created.

PC DOS 6.3's E Editor also comes with a set of powerful macros. <DRAW> allows you to do freehand drawing with the arrow keys; <SORT> lets you mark a block of data and sort it; <ADD> lets you do column addition; and <BOX> lets you draw boxes of arbitrary size.

Another powerful, new feature in PC DOS 6.3's E Editor is the ALL command. The ALL command creates a file



called *.ALL, which contains all instances of a search. Once this file is created, you can use the key combination Ctrl+Q to toggle between the *.ALL file and the actual file where the occurrence took place. This feature is very useful when debugging code.

System Management Tools

Boot Options. PC DOS 6.3 offers three options at bootup: Clean Boot, Interactive Boot, and Selective Boot.

Clean Boot lets you bypass loading the CONFIG.SYS and AUTOEXEC.BAT files, and sends you directly to the C:> prompt. If you have ever been unable to boot DOS because of errors in your CONFIG.SYS file, clean boot will save your day!

Suppose, for example, a new device driver was added to your system, and modifications were made to the CONFIG.SYS file. The next time you reboot, your computer hangs. It's time for clean boot! Reboot the system and press F5 when you see the DOS start-up screen. Your system will bypass its startup files and go to a C:> prompt. At that point, you can edit your CONFIG.SYS and correct the problem. (In previous versions of IBM DOS, and in other versions of DOS, you would have to boot from a bootable DOS diskette to get a C:> prompt.)

To bypass the DBLSPACE.BIN (compression driver) file in addition to CONFIG.SYS and AUTOEXEC.BAT, press Ctrl+F5. (This feature is new in PC DOS 6.3.)

Interactive Boot lets you step through your CONFIG.SYS and AUTOEXEC.BAT (new in PC DOS 6.3) files. At the DOS start-up screen, when you press F8, the system will step through these files line-by-line, prompting you to enter Y or N each time.

For example, suppose you want to use a large spreadsheet, and suppose you also typically activate your LAN system when you boot. However, this time the spreadsheet is so large that you run out of memory – your system cannot handle both the large spreadsheet and your LAN. Reboot your computer and press F8. When you are prompted to “load LAN device driver,” reply with N. Now you have plenty of memory for your spreadsheet. (In previous versions of IBM DOS, or in other versions of DOS, you would have to edit your CONFIG.SYS, REM out the LAN device drivers, and reboot. When done, you would have to go back to your CONFIG.SYS and remove the REM statements.)

Lazy writing adds significant performance improvements to write operations.

The key combination Ctrl+F8 at the DOS start-up screen bypasses loading the DBLSPACE.BIN file in addition to stepping through the CONFIG.SYS and AUTOEXEC.BAT files.

Selective Boot lets you set up multiple configurations within a single computer. At bootup, a customized menu is displayed, indicating the different configurations the computer can boot into. This feature provides a lot of flexibility when managing multiple configurations, especially in office environments where one computer might have various functions.

Selective boot is done within your CONFIG.SYS and AUTOEXEC.BAT by placing groups of statements into *configuration blocks*. When the computer is

booted, PC DOS displays a menu and asks you to choose the configuration you would like to run.

As an example, suppose your system is used to dock a laptop, link to the office LAN, and test memory-intensive games. You run out of memory when you are linked to the LAN and you attempt to dock a laptop or run a game.

Within your CONFIG.SYS, set up three configuration blocks: block 1 only loads the docking drivers, block 2 loads the LAN drivers, and block 3 loads no drivers. Now, when you boot your system, a menu appears prompting you to select the configuration you require at that time.

In previous versions of IBM DOS, or in other versions of DOS, you would have to edit your CONFIG.SYS and remove or add the appropriate drivers each time you wanted to do one of the tasks, and then reboot your system each time.

Improved SMARTDRV. PC DOS 6.3 has an improved SMARTDRV disk cache, increasing disk performance. In PC DOS 6.3, the SMARTDRV capability has also been given to CD-ROMs.

SMARTDRV now includes lazy writing (delayed writing), which stores information to be written to your hard disk, and writes it to the disk later when the system is less busy. Lazy writing adds significant performance improvements to write operations.

Also, SMARTDRV is now an executable (*.EXE), allowing you to change, enable, or disable caching dynamically. In previous versions of IBM DOS, or in other versions of DOS, you would have to modify the CONFIG.SYS file and then reboot your system.

Defragmentation reorganizes files on your hard disk and arranges these files in a more consistent manner. This feature minimizes the time it takes your

computer to retrieve data. In PC DOS 6.3, the defrager now uses extended memory to defrag even large hard disks efficiently and safely.

Due to the way the FAT file system is structured, fragmentation happens to all disk drives over time. As more data is added to or replaced on your disk, the fragmentation increases. By consistently running Defrager, your hard disk will remain optimized.

Defragmentation of large disks is time-consuming. With PC DOS's scheduler program, you can schedule the defrager to run during times when the computer is not being used. (See "Program Scheduler" below.)

NEW! Interactive Batch Processing brings the step-through feature found in interactive boot to batch files. Its function is similar to interactive boot, except that you do not have to reboot your system in order to step through a batch file.

A new switch (/Y) has been added to the COMMAND command to enable you to step through a batch file. The syntax for doing this is:

```
command /y /c filename.bat
```

With the /Y switch, you can step through the batch file one line at a time, making it easier to troubleshoot batch files.

Example of use: Similar to the example in the "Interactive Boot" section.

CHOICE is a handy command for getting the user's input from batch files. Using CHOICE in a batch file, you can specify what to prompt the user. The user types in a response that is passed back in an error-level return code to your batch file. Now you can do conditional branches on the return code within your batch file.

DELTREE is a command that lets you delete an entire subtree or hierarchical

directory structure of files at once. Now, if you want to delete a directory, you no longer need to delete all of its files and subdirectories first. Instead, use the DELTREE command followed by the name of the directory you want to remove. PC DOS 6.3 then deletes all subdirectories and files below the level of the directory you deleted.

For example, suppose you have a directory called WORDPERF, and within that directory are eight subdirectories. You want to remove the WORDPERF directory and its directory hierarchy. Simply type DELTREE WORDPERF to remove all files, subdirectories, and the WORDPERF directory itself.

Contrast this simple command with the steps you would have to take in previous versions of IBM DOS or in other versions of DOS. You would have to go into each subdirectory, delete the files, back out one level, and remove the subdirectory. In the WORDPERF example, this process would be repeated eight times, one for each subdirectory. Then, once all subdirectories are removed, you would have to back out another level and remove the WORDPERF directory.

MOVE is a command that lets you move a file or group of files from one directory to another, or move a whole directory hierarchy from one directory to another directory. The files in the original (source) directory are removed and placed in the new location. Finally, the source directory is itself removed.

For example, if you want to move files that reside in C:\DOCUMENT to E:\DOCUMENT, simply type

```
MOVE C:\DOCUMENT\*.* E:\DOCUMENT
```

In previous versions of IBM DOS, or in other versions of DOS, you would have to copy the files from C:\DOCUMENT to E:\DOCUMENT, then go back to C:\DOCUMENT and delete the files, then back out one level and remove the directory.

NEW! Smart Upgrade. PC DOS 6.3 installation has been modified to understand the multi-configuration support delivered in PC DOS 6.1 and MS-DOS 6.0/6.2. Previously, when upgrading or changing to any of these versions of DOS, the installation process would not recognize multi-configuration support, and installation files would be randomly placed within the CONFIG.SYS file. Now, in PC DOS 6.3, installation is aware of the multi-configuration, and files are added to the CONFIG.SYS correctly.

NEW! CD-ROM Support (MSCDEX 2.23). PC DOS 6.3 provides support for accessing CD-ROM drives by using MSCDEX.

NEW! Thousands Separator. PC DOS 6.3 now displays a separator for thousands digits. This function displays specific country versions of separators, and it works with the DIR, DSK, FORMAT, and MEM commands.

A switch has been provided to enable the user to turn this feature off. This switch is vital for staying compatible with existing programs that read data from these commands and do not expect a separator character.

NEW! No-Swap DISKCOPY. PC DOS 6.3's no-swap DISKCOPY copies the entire contents of one diskette to another by using the system's hard disk as temporary storage. This feature eliminates the task of swapping diskettes multiple times when using DISKCOPY.

NEW! File Overwrite. When using the COPY, XCOPY, or MOVE command, PC DOS 6.3's file overwrite feature prompts the user before overwriting files that have identical names. This feature helps eliminate the accidental deletion of files. PC DOS 6.3 has a switch that allows the user to turn this feature off.

Data Recovery

Undelete is a full-function, full-screen undelete program that allows you to re-access files that have been inadvertently deleted. Both DOS and Windows interfaces are provided.

PC DOS 6.3's undelete supports three levels of protection. The lowest level of protection, *Standard DOS*, gives you the ability to recover files that have not been overwritten.

The second level of protection is *Tracker*. When you are using Tracker and you delete a file, DOS first records the address of each cluster of the file, then makes those clusters available for use by other files. If no other file has overwritten those clusters, you have a very good chance of recovering the deleted file.

The highest level of protection is *Sentry*. Files protected by Sentry are saved to a hidden directory and can be retrieved by Sentry in perfect condition. If Sentry detects that your hard disk is running out of space, it will automatically remove the oldest files that have been stored in the hidden directory, freeing up space for newer ones. This feature can be customized by the user.

As an example of Sentry: You have inadvertently deleted a very important spreadsheet. By invoking PC DOS 6.3's undelete program, a list of previously deleted files is displayed. Click on the file you want to undelete, and press F8. Your important file is recovered.

Undelete File Viewers let you examine the contents of files prior to undeleting them. Files are shown in their native format when it can be determined; otherwise, they are shown in either text or binary format. Windows file viewers are available when using the Windows Undelete interface.

For example, the file you inadvertently deleted is listed with hundreds of other files. You don't know the name of the file, but you know its contents. Click on any of the files listed, and press F4. That file now appears in its original format. Continue through the list until you find the correct file. Once you find it, press F8 to undelete it. (In previous versions of IBM DOS, or in other versions of DOS, you would have to undelete each file and bring up the file application program to view it. Once found, you then had to go back and delete each file that was undeleted.)

Undelete is a full-function, full-screen undelete program that allows you to re-access files that have been inadvertently deleted.

Program Scheduler (Central Point)

PC DOS 6.3 includes a full-functioning program scheduler that allows you to run any program or DOS command automatically. Having tools such as BACKUP, DEFRAG, and AntiVirus is great, but having to sit at your desk and wait for them to run is not. With PC DOS 6.3's program scheduler, you can schedule all those utility programs and any other DOS programs or commands to run at a certain time. PC DOS 6.3's scheduler comes with an easy-to-use calendar interface.

PCMCIA Support (Phoenix)

PCMCIA** technology is being incorporated into many systems being designed by IBM and other OEMs. This new technology offers flexibility, portability, ex-

tendibility, and ease of use. Flexibility means the cards work in all types of personal computers. Portability denotes cards that are small, consume very little power, are very rugged, and can be inserted and removed easily. Extendibility means that different cards can be plugged into the same socket, and through a process called hot insertion/removal, cards can be inserted and removed while the power is on.

The architecture of PCMCIA starts with the PC card. These cards are the size of a credit card and provide support for memory or I/O. A PCMCIA card is inserted into a slot. The socket is attached to an adapter, which is a controller that connects one or more PCMCIA sockets to a system bus.

Socket Services is software that talks to the adapter. It is at the BIOS level, which is low-level software that provides a universal interface for the hardware. Socket Services is a DOS-loadable driver, and can either be loaded in the CONFIG.SYS file or run as a TSR from DOS.

Socket Services works with any OEM 2.0 level socket services in any OEM or IBM system. It works with existing PC cards that meet the 2.0 specification. Socket Services complies with the PCMCIA Card Services 2.0 interface specification as defined by the PCMCIA.

Card Services is operating-system code that provides a standard application programming interface (API) for clients. PC DOS 6.3 has incorporated PCMCIA Card Services at the 2.0 level.

Card Services interfaces directly with Socket Services, and is responsible for coordinating all access to the PC cards and allocating system resources among Card Services client drivers.

Card Services has many responsibilities. One is complete management of all system resources available for PC cards.

Another is its responsibility for managing client drivers that are written for a specific PC card. These drivers register with Card Services at initialization. Card Services then provides information to these drivers as PC card events occur. Card Services is also responsible for handling all PC card hardware events, such as card insertion, extraction, and status changes, battery flow, and card ready/lock.

When a hardware event occurs, Card Services issues a callback to any registered client to inform it that a hardware status change has occurred. Card Services has complete control of power management. When a suspend request occurs, Card Services puts each PC card into a low-power mode (if it is supported by the PC card). Card Services can do this because it is aware of Advanced Power Management (APM), and is connected to the system to broadcast this function. At the resume request, Card Services determines if the PC card has been changed or removed and, if so, the PC card is reconfigured as if it has just been inserted.

Device Drivers and Utilities. Eight device drivers and utilities are now discussed.

Super Client Driver is a PC DOS 6.3 loadable driver that becomes a Card Services client driver. This driver is designed to support the configuration of a wide array of PC cards.

The Super Client Driver is a collection of client drivers that perform PC card configuration. After configuration, the PC card operates exactly as though it were a permanent component when the system was started.

The set of client drivers in the Super Client Driver is determined at link time. Each individual client driver can either be linked separately or with a group of other client drivers.



The Super Client Driver does not support every PC card in the industry. If the Super Client Driver does not support a PC card, then a client driver to support the card will have to be provided by a third-party supplier.

Information Utility Program provides status information about each PC card that has been installed into the PC card socket. This utility has DOS and Windows interfaces.

When the information utility is run, the status of each socket is displayed, advising the user whether a socket is empty or if it is in the process of configuring a card. If a card has been installed, the utility displays the name and information about the PC card. In the case of a non-configurable card, a message is displayed informing the user of that status.

Super Memory Technology Driver (MTD) is a DOS driver that has been designed to support read, write, erase, and copy functions for PC memory cards. A memory technology driver (MTD) is needed for each specific memory technology that requires read, write, erase, and copy functions. The Super Memory Technology Driver is similar in design to the Super Client Driver, and it consists of a collection of MTDs that operate only with Card Services.

Virtual Driver for FAT Block Devices supports ATA-compatible fixed disks and SRAM cards formatted in a FAT-structured format.

Windows VxD Driver enables fax and modem PC cards to be fully operable under all Windows sessions when inserted into a PCMCIA socket. Card configuration is always performed under Windows. However, Windows allows the

fax/modem cards to be available only to the current Windows program. The VxD driver resolves this restriction and makes the fax/modem cards available to every Windows session.

SRAM Format Utility is a formatting program for SRAM PC cards that are accessed as either drive A or drive B. This utility is necessary because a DOS format program limits the capacities available for drives A and B. SRAM cards can also be made bootable using this utility program.

FAT Diskette Emulation. A driver is provided for FAT diskette emulation. When this driver loads, it performs diskette drive emulation on the specified socket, and then registers it with Card Services as a memory client.

PC COPY and PC ERASE for FLASH PC Cards are utilities designed especially for FLASH PC Cards. These utilities are not included in PC DOS 6.3, but are available for OEM customers on an as-needed basis.

Pen Support (IBM/CIC)

Another PC DOS 6.3 advantage is its support for pen hardware. PenDOS, included in PC DOS 6.3, gives PC hardware manufacturers the ability to include pen enablement at no additional cost.

PC DOS 6.3's PenDOS supports existing unmodified mouse-aware DOS applications as well as "pen-aware" applications.

The pen extensions offered in PC DOS 6.3 are mouse emulation, gesture recognition, numeric recognition, a writing window, and a pop-up soft keyboard.

Mouse emulation lets any current existing DOS "mouse-aware" application

function with a pen the same way it would use a mouse.

Gesture recognition brings common editing gestures to your current DOS applications if they support the gestures: undo, cut, paste, and so on. For example, if you are using a word processing program for DOS that is mouse-aware, you can use the pen to bring down menus, select items, and even highlight text, and then make an 'X' gesture to delete the highlighted text.

With both mouse emulation and gesture recognition, PC DOS 6.3 makes current DOS applications function easier, and it makes the application more intuitive to use.

Numeric recognition lets pen-aware applications or standard DOS applications accept numeric handwritten text. With a pen-aware application, numbers can be hand-written directly into fields and recognized. In standard DOS applications, a writing window would need to be displayed. Once displayed, numbers could then be hand-written, recognized, and then sent to the DOS application.

The pop-up soft keyboard is a feature that displays a keyboard on the screen which can be used to send keystrokes to any DOS application.

The one pen extension not shipped in PC DOS 6.3 is alphanumeric handwriting recognition. A coupon has been included in the PC DOS 6.3 User's Guide that allows PC DOS 6.3 users to upgrade to the full-featured PenDOS operating system for \$39.95, affording a \$40.00 savings.

Performance

PC DOS 6.3 has gone through extensive hours of optimization and "scrubbing." Some of these enhancements are in the

area of 386/486 code optimization, improved interrupt handling, command-processor tuning, and character device handling.

Performance tuning is a slow, resource-intensive activity. Starting with PC DOS 6.1 and continuing with PC DOS 6.3, IBM has an ongoing process for continuing to enhance PC DOS with the performance characteristics that customers are demanding.

Robert Rapuano is an advisory programmer in Entry Operating Systems Product Planning within IBM's Personal Software Products division. He is responsible for PC DOS product planning and strategy. Also during his IBM career Robert was manager of OS/2 performance evaluation and a programmer developing OS/2 and tools. He can be reached via internet at rapuano@bcrvml.vnet.ibm.com.

Nestor Miranda is a PC DOS Product Planner in IBM Personal Software Products Division, Boca Raton, Florida. In addition to product planning, he has been instrumental in building product awareness for the PC DOS product through writing publications, providing media support, and attending business shows. He joined the IBM account in 1992 as a lead in the Distributed Application Environment test team, then joined the Mobile, Voice, and Pen organization as a planner for announcing PenDOS 2.2. Nestor has a BS degree in computer science from Barry University and is working toward his MS in management there. In his spare time, he manages and plays third base on a semi-pro baseball team in Homestead, Florida.

PC DOS 6.x Disk Compression and Defragmentation

Data compression is not new – the mathematical theory, generic compression algorithms, and DOS-based programs have been around for many years. But compression of entire disks in a manner virtually transparent to the user is a recent innovation, and is certainly one of the most significant enhancements to IBM PC DOS 6.x.

This article discusses data compression techniques; file compression; disk compression; sectors, clusters, and slack space; fragmentation and DEFRAG; installation; and utilities.


This article is excerpted and adapted from Upgrading to PC DOS 6.1, one of a series of Red Books published by the IBM International Technical Support Organization (ITSO) in Boca Raton, Florida. The IBM order number for this publication is GG24-4133. Its authors are:

*Heinz Schelter-Kaufmann, IBM
Deutschland GmbH
Dwight Ronquest, IBM South Africa
Laurie Rose, IBM United Kingdom Ltd.
Tim Sennitt, IBM International
Technical Support Organization,
Boca Raton, Florida, who was also
the advisor for the project.*

Data Compression Techniques

The following sections describe some of the most popular methods of compressing data, and give some background information about the various compression theories.

Run-Length Encoding: The simplest method of data compression involves scanning for repeated occurrences of particular characters, and replacing them



with just one instance of that character and a count of the number of times to repeat it. An *escape* flag is also needed to inform the decompressing program that this is a compressed sequence.

This method works best on plain text, which often contains long sequences of blanks, dots or underscores. It also works well on some graphics or image files where large areas of a particular shade are represented by consecutive color attributes. Some graphics file types already implement this type of compression in their internal format.

Run-length encoding is less effective for executable code, though a program often contains text (messages), uninitialized arrays, and data buffers that compress quite well.

Huffman Compaction: Computers store characters in 8-bit bytes; each byte can be a value between 0 and 255. The ASCII codepage assigns every character (a-z, A-Z, numbers, punctuation, and graphics characters) to a particular value. Each character therefore takes up exactly 8 bits of disk space. Although this is very convenient for organizing and processing data, there are denser ways of storing it.

If we only wanted to store the letters A-Z, we could do so in just 5 bits (enough to handle 32 characters). Six bits would allow us to encode both upper- and lower-case letters and also numbers. Seven bits is sufficient for the basic 128-character ASCII set; indeed, 7-bit ASCII is sometimes used in data transmission. But, because these schemes cannot

<sp> - 3929	R - 951	H - 497	V - 223	Q - 27
E - 1839	A - 898	P - 423	W - 218	Z - 4
T - 1273	N - 834	L - 413	G - 198	J - 3
S - 1183	D - 715	U - 388	B - 192	
O - 1116	M - 612	F - 339	K - 127	
I - 961	C - 577	Y - 261	X - 87	

Figure 1. Distribution of Letters in PC DOS 6.1 README.TXT

<sp> - 000	O - 11100	D - 10101
E - 001	I - 10001	M - 10110
T - 010	R - 10010	C - 10111
S - 011	A - 10011	
	N - 10100	
H,P,L, etc - '110' + 4 bits		
All others - '111' + the character's 8 bits		

Figure 2. Huffman Compression Using Rules Specified

handle all possible characters, they are not really suitable for storing programs or word-processed documents on disk.

Huffman compaction uses a variable number of bits to encode characters – some are stored in fewer than 8 bits, while others take more. The more frequently a character occurs, the shorter its bit-string (just as in Morse Code). Each bit pattern must be unique so that the decompressing program knows where each character starts and ends within the 8-bit byte.

“Real-life” Huffman compaction programs may be adaptive – that is, they scan the source data and analyze the distribution range of characters, and choose a packing methodology most suited to that data. We will look at a simplified example to illustrate the technique, but this is not how it really works!

You probably have a feel for the typical distribution of letters in the English language, but to quantify it better, we wrote a program to count the letters in a typical text file (which happens to be the PC DOS 6.1 README.TXT). Ignoring case, the result is shown in Figure 1.

This pattern is probably typical for English, and should come as no surprise to people who are familiar with the board game Scrabble**.

We could perform Huffman compression on text using the following rules:

1. If the character is among the top four in frequency, set the first bit to 0, then use two more bits to identify which of the four. This rule requires 3 bits total.

For rules 2, 3, and 4, the first bit is set to 1.

2. If the character is among the next eight, set the second bit to 0, then use three more bits to identify which of the eight. This rule requires 5 bits total.

For rules 3 and 4, the second bit is set to 1.

3. If the character is among the next 16, set the third bit to 0, then use four bits to identify which of the 16. This rule requires 7 bits total.

For rule 4, the third bit is set to 1.

4. Otherwise, use the character's 8 bits as they are. This rule requires 11 bits total.

The optimum algorithm clearly depends on the slope of letter distribution.

Using the rules above, Huffman compression would generate the bit patterns shown in Figure 2.

As another example that uses these rules, take the phrase ITSO RED BOOK which has 13 characters that would normally occupy 104 bits. Five-bit ASCII code would require 65 bits. But it can be compressed into just 60 bits using the encoding pattern shown in Figure 3.

This technique works only when the distribution of byte values is non-random. It is quite effective for plain text, giving a compression factor of around 30 percent. It works reasonably well on executable code too, since certain instructions tend to be more common than others, and (as already mentioned) programs usually contain message text and buffer space.

I	T	S	O	<sp>	R	E	D	<sp>	B	O	O	K
10001	010	011	10000	000	10010	001	10101	000	1111010	10000	10000	1111011
10001010	01110000	00010010	00110101	00011110	1010000	10000111	1011...					
Hex: 8A	70	12	35	1E	60	87	Bx					

Figure 3. Huffman Compression of the Phrase “ITSO Red Book”

Lempel-Ziv-Welch Compression: In a normal text file, common words and word fragments (for example, THE, ING, ION, COM, ER, TO) appear many times. In 1977, Lempel and Ziv developed a technique that analyzes a file for recurring phrases. These phrases are put into a data dictionary. After that, each recurrence in the text is replaced by a *token* that represents the original character string. This method, later enhanced by Welch, is also known as LZ or LZW compression.

An example may make it clearer, although the power of the technique may not be immediately apparent with such a small piece of text. Take the phrase:

THESE ARE THE VOYAGES OF THE
STARSHIP ENTERPRISE.

We could build the following dictionary:

1 THE
2 SE
3 AR

after which our phrase becomes:

12 3E 1 VOYAGES OF 1 ST3SHIP
ENTERPRI2.

Clearly, the longer the text, the greater the likelihood of more common strings appearing, and the higher the overall compression.

Implementation: The LZW method (with further Huffman-like compaction of the data directory) forms the basis of most compression products available today. Nevertheless, there can be a considerable difference in the elegance and skill of the implementation, which ultimately determines the product's speed, compression ratio, memory overhead, robustness and general functionality. Of course, fast, reliable decompression is equally important.

File Compression

The above methods illustrate some ways that raw data streams can be compressed. What the end-user needs is a compression program that operates on

normal files. In the PC world, these programs tend to be released as Shareware or Freeware products. Examples include PKZIP**, LHARC, LZEXE, and ZOO.

File-compression programs are commonly used to pack files into a convenient format for transmission, such as a bulletin-board download or a mainframe-to-PC file transfer. Clearly, the smaller the data package, the sooner the transfer completes, saving both time and line connection costs.

To make compression really effective, it needs to be an integral part of the DOS file system.

For convenience, the compression programs can usually combine several source files into one compressed output file, often known as an *archive*. Only one file transfer operation is necessary. Then the decompression program separates the archive back to its original component files. The original subdirectory structure can also be preserved.

Archives can be made self-expanding by attaching an executable piece of header-code. This means the recipient of the file does not need a separate decompression program – the recipient just “runs” the archive file.

Today, many software products are manufactured in compressed form to reduce the number of diskettes in the package. Their compressed files can sometimes be recognized by an unusual character in the file extension (such as a dollar or underscore).

Manual file compression can certainly be used to increase effective disk stor-

age, if you remember to build archives, decompress them when needed, and re-compress afterwards. This process could be automated to a certain extent using appropriate batch files.

This approach is taken a stage further in a self-expanding, self-executing program. Many of the EXE and COM files in PC DOS 6.x are internally compressed using PKLITE**. The main body of the file is stored in a compressed form on the disk (whether disk compression is installed or not). When the program is executed, its uncompressed PKLITE header code runs first. This header code loads and decompresses the main program into memory, then passes control to it. The result is a considerable saving of disk space and load time.

However, this scheme can make problem determination more difficult, for the following reasons:

- You cannot easily patch or edit the program.
- A text-scan utility won't find the originator of an error message.
- Tracing or disassembling the program (for example, with DEBUG) is awkward.

Also, since the file is already compressed, DOS disk compression will yield less overall benefit than might be expected.

Ideally, compression should not be the responsibility of the PC user, or of individual programs. To make compression really effective, it needs to be an integral part of the DOS file system.

Disk Compression Schemes

Today's disk compression products (such as SuperStor/DS, Stacker**, and DoubleSpace) all operate at and below the DOS file-system level. This means they can be completely transparent to the user. They all work in a broadly similar way.

PC DOS with Stacker 3.0			PC DOS with SuperStor/DS		
IBMBIO	COM	40694	IBMBIO	COM	40646
IBMDOS	COM	37174	IBMDOS	COM	37174
STACKER	<DIR>		COMMAND	COM	52797
DOS	<DIR>		DBLSPACE	BIN	50436
STACVOL	DSK	49960960	DBLSPACE	INI	42
COMMAND	COM	52589	DBLSPACE	000	49575936
STACKER	LOG	3103			
AUTOEXEC	BAT	399			
CONFIG	SYS	352			

Figure 4. Host Drive after Compression Is Installed

How Disk Compression Works: Here is an outline of the steps in compression.

1. During installation, a hidden file (the *compressed drive file*) is created, and all existing files are moved and compressed into it. Some free space – a few MBytes – must be available before the product is first installed, because it needs some elbow room to begin the compression process.

As installation proceeds, the compressed drive file grows into the freed-up space until it occupies most of the physical partition (often called the *host or uncompressed drive*). The DOS kernel files, IBMBIO.COM and IBMDOS.COM, and the disk compression driver itself stay uncompressed on the host drive, because they must be available at boot time.

2. The compression driver maps the compressed drive file to a DOS drive letter. From now on, it behaves just like a normal DOS disk.
3. Drive letters are adjusted so that the compressed drive appears as C: (then D: and onward if more than one drive is compressed). The physical host drive is mapped to a higher drive letter, and remains accessible if necessary.
 - In their DOS 6 implementations, DoubleSpace and SuperStor/DS are fully integrated into the base operating system. The driver is loaded by, and becomes part of, the DOS I/O kernel early in the

boot process. CONFIG.SYS resides in the compressed volume, because compression support is already fully operational before CONFIG.SYS is processed.

This also means that all DOS and other device drivers can reside on the compressed drive. For the user, this greatly simplifies management of the hard disk.

- Stacker (to v3.0) loads last as a CONFIG.SYS device driver, so CONFIG.SYS itself, some DOS, and user device drivers must reside on the unstacked (uncompressed) drive. Stacker has to check for synchronization of the stacked and host drive contents.
 - Stacker 3.1 exploits the hooks in the DOS 6.0 kernel to load seamlessly in the same way as DoubleSpace and Superstor/DS.
4. All I/O data transfers to the compressed drive pass through the compression/decompression code. This is completely transparent to programs and the user. But note that programs which bypass DOS or BIOS interfaces may fail, and could corrupt the compressed drive file.
 5. After compression, disk space is nominally doubled. The compression driver handles calls that query the disk size and free space (as used by DIR, CHKDSK, and application installation routines) and reports back values that are double the original

values. DIR reports the real (uncompressed) size for individual files; only the free-space value is doubled.

Figure 4 shows how the host drive might look after compression is installed. Normally, you never need to bother about the host-drive contents.

Differences: In preparing this article, we did not run any formal comparisons of SuperStor/DS, DoubleSpace, or Stacker. Speed and compression ratios are reputedly very similar. However, there are other factors that should be considered when comparing disk-compression products:

- What support utilities are provided? Can they:
 - Report the compressed volume status and validate data integrity?
 - Recover damaged files?
 - Measure the actual compression ratios, and adjust the free-space factor accordingly?
 - Alter the compressed and host disk size?
 - Easily uninstall the product altogether?
- How well is error recovery handled?
- How much memory does the driver use? What kind?
- How compatible is it with other real-time applications?
- How does it work in a LAN environment?
- Can you adjust:
 - Cluster size?
 - Compression versus speed-tuning?
- Most important: How reliable is it?

Compatibility: Stacker, Superstor/DS, and DoubleSpace use different compression techniques and formats, so you cannot change from one product to the other without decompressing the drive first.

However, PC DOS 6.x comes with SuperStor/DS, a completely different product that *can* access DoubleSpace-format disks. (Indeed, SuperStor/DS cannot access traditional DR DOS 6.0 SuperStor disks). You can therefore easily upgrade from MS-DOS 6.0 to IBM PC DOS 6.x.

Of these compression products, only Stacker currently has an OS/2 version. This version uses the same format, so a stacked (compressed) disk is accessible from both DOS and OS/2.

Compression Ratios

In general, you will be lucky to double your disk space. Although the compression product will report twice the space, a more reasonable expectation might be 1.5 to 1.7 times. This disparity can cause anomalies when the disk is nearly full.

In practice, much depends on the type of files on your disk. As we've already seen, plain text and some graphics files compress extremely well. Executable code compresses reasonably. The worst candidate is a file that is already compressed. You can be reasonably certain a file is already compressed if it has a tell-tale filetype (e.g., .ZIP, .ARC, .ZOO, .LHA), but you may not know whether executables are packed by PKLITE or something similar.

Also, many recent PC games (particularly those supplied on more than three or four diskettes) install their files in a compressed form, further reducing the scope for real disk-doubling.

Some disk-compression products recognize that a file is already compressed, and they store it directly. This avoids the compression/decompression overhead when the file is accessed. Of course, if you want to create an archive for file transfer, you still need to run PKZIP or a similar utility.

DIR: The actual compression ratio for a file can be listed by using a new DIR

4201	CPI	26420	06-29-93	12:00p	2.1 to 1.0
4208	CPI	115	06-29-93	12:00p	16.0 to 1.0
ADMIN	PRF	715	06-29-93	12:00p	16.0 to 1.0
ANSI	SYS	9029	06-29-93	12:00p	2.1 to 1.0
APPEND	EXE	7728	06-29-93	12:00p	1.0 to 1.0
ASSIGN	COM	5022	06-29-93	12:00p	1.0 to 1.0
ATTRIB	EXE	10095	06-29-93	12:00p	1.0 to 1.0
AUTORUN	PRF	215	07-06-93	2:18p	16.0 to 1.0
CHECKUP	DB	7774	07-06-93	2:18p	1.8 to 1.0
CHKDSK	COM	12633	06-29-93	12:00p	1.0 to 1.0
CHOICE	COM	1586	06-29-93	12:00p	4.0 to 1.0
CMOSCLK	SYS	855	06-29-93	12:00p	8.0 to 1.0
COMMAND	COM	52589	06-29-93	12:00p	1.4 to 1.0
COMP	COM	8561	06-29-93	12:00p	1.9 to 1.0
CONTACT	LST	615	06-29-93	12:00p	16.0 to 1.0
COUNTRY	SYS	18142	06-29-93	12:00p	3.7 to 1.0
CPBACKUP	EXE	409646	06-29-93	12:00p	1.6 to 1.0

Figure 5. Sample DIR /C Output

parameter. Type DIR <filename> /C to see the results. DIR /OC lists matching files sorted by compression ratio. The C parameter may, of course, be combined with other DIR options. Figure 5 gives an example – the start of the DOS directory. Notice the wide range of compression ratios.

Slack Space: All disks and diskettes are physically formatted in 512-byte sectors. DOS uses space in quantities called *clusters* (or *allocation units*), where a cluster is one or more adjacent sectors on the disk surface.

DOS keeps track of cluster usage in the volume's File Allocation Table (FAT, which gives the DOS file system its name). The FAT contains one entry for each cluster; the entry itself is either a chain pointer to the next cluster in the file, or zero indicating that the cluster is free, or a flag marking it as defective.

Over time, as files are created, erased, expanded, and shrunk, the chains of clusters making up files can become very fragmented. A single file might become scattered all over the disk surface. This affects performance adversely, since much more disk activity (head-seeks and rotational delay) is needed to load the whole file. It also loses some of the benefits of caching; a read-ahead cache

assumes that physically adjacent sectors belong to the same file, which won't be true if the file is fragmented.

A sector has 512 bytes. What is the best cluster size?

At one extreme, a cluster can be just one sector. This is indeed the case for HD diskettes. But what about hard disks?

- A cluster of just one sector would make the FAT very big. Consider, for example, a 400 MB hard disk that has some 800,000 sectors. Each FAT entry is 16 bits (2 bytes), so the FAT alone would occupy 1.6 MB (plus, there is a backup copy on the disk for safety). This size is too large for DOS to maintain and search efficiently.
- As space is allocated and freed in lots of little pieces, the disk would become hopelessly fragmented very quickly.

The solution is to use larger clusters. DOS selects a cluster size of 2 KB, 4 KB, 8 KB, or more when formatting a hard disk, depending on its capacity. These cluster sizes make the FAT a manageable size, and help to reduce the degree of fragmentation somewhat. They also increase the chance of related clusters being in the disk cache.

So, why not choose a very large cluster size? The problem is that the whole cluster is allocated to the file, even if it doesn't need all of it. Consider a disk formatted with 8 KB clusters. A one-byte file still occupies 8 KB of disk space. A file of 8193 bytes occupies 16 KB of disk space, and so on. The same applies to subdirectory entries as well.

On average, half a cluster – up to 4 KB – is wasted for *every* file and subdirectory on the disk. This *slack space* can account for a significant amount of the entire disk's capacity. For example, if the average file size is 40 KB (not untypical) on our 400 MB disk, then 10 percent of its capacity could be wasted. It will be much worse if there are lots of small files. So, in summary, the cluster size is a delicate compromise chosen to balance acceptable performance without wasting too much disk space.

Compression and Slack Space: As we saw earlier, disk compression files occupy most of their host volume, so there is virtually no slack space. Although the compressed disk file behaves like a FAT drive, disk-compression programs use their own advanced, proprietary internal file- and space-allocation techniques to manage the contents of the compressed volume file.

One consequence is that Stacker and SuperStor/DS (but not DoubleSpace) can use cluster sizes smaller than those of the host drive, with much less impact on performance. Indeed, SuperStor/DS can run with a cluster size of only 512 bytes. This size dramatically reduces wasted slack space, and can make a significant contribution to overall disk-doubling, even though individual file compression ratios don't achieve the 2:1 target.

Performance

One might expect that the overhead of compressing and decompressing data in real-time would take its toll on system

performance. In practice, much depends on the system hardware.

As a rough guide, disk compression on a 20 MHz 80386SX system probably won't cause any noticeable degradation of performance. The overhead of the increased compression-code pathlength is offset by the reduced volume of data transferred to or from the hard disk. A computer with a slower processor may experience reduced performance, whereas one with a fast processor may actually run faster. On the other hand, computers with slow processors tend to have slower disks, so there may still be an overall balance.

Bootting from Diskette: With some disk compression schemes, it may not be possible to access the compressed hard disk if you boot from diskette, since the driver is not loaded. In PC DOS 6.x, FORMAT /S copies the compression driver to diskette, and the kernel loads it automatically, so the compressed disk is always available. This should mean that DOS Corrective Service can be applied as normal – simply boot the corrective-service diskette.

Compressed data may not be accessible if you boot a different operating system, and care should be taken not to damage the compressed drive file in such circumstances. Be very careful running any disk maintenance or recovery tools this way.

DEFRAG

As explained earlier, the cluster method of space allocation in the FAT architecture inevitably results in eventual file fragmentation.

The DEFRAG utility in PC DOS 6.x should be run from time to time. It reorganizes files and directories so that they occupy contiguous clusters on the disk, thereby eliminating fragmentation. When defragmentation is complete, all files are contiguous, and disk performance can increase noticeably.

DEFRAG itself measures the degree of disk fragmentation, and advises whether it is worth running, and (if yes) which method is recommended.

Fragmentation occurs in both compressed and uncompressed disks.

CHKDSK should be run before DEFRAG. CHKDSK can also be used to determine which files are fragmented. The format of the command is `chkdsk <filename>`.

Figure 6 illustrates the DEFRAG display screen.

There are two levels of DEFRAG optimization:

- Full, which defragments all files and directories, and also ensures that all free space is contiguous.
- Files only, which defragments all files and directories, but may leave areas of free space between them.

Moving subdirectory entries to the top of their parent directory helps DOS performance when locating files deep within the disk's tree structure.

DEFRAG can also sort files (by name, date, size, etc.) when relocating them on the disk. This can make raw DIR listings neater, although DIR itself now supports sort parameters. DIR preferences can also be permanently configured by setting the DIRCMD environment variable. Any DIRCMD sort override is always in effect, whereas DEFRAG sorting becomes lost over time.

Installation of PC DOS 6.x Compression

Important! It is vital that you read the documentation that is shipped with PC DOS 6.x compression before undertaking installation.

Before Installation: You can use only one compression program on your computer at one time. Therefore, it is very

important that you either remove any existing incompatible compression from your disk, or continue to use it without installing the PC DOS 6.x compression.

The SuperStor/DS compression in PC DOS 6.x is currently compatible only with Microsoft's DoubleSpace compression. SuperStor/DS is not currently compatible with any other compression program.

To be able to recover if disaster strikes, make sure that you create a bootable diskette after installing the compression software. For more information, see the later section "Recovery Diskette."

Back Up Your System: You must choose a backup method that will be compatible and can be restored to PC DOS 6.x when it is installed. You may wish to install CPBACKUP, which is part of PC DOS 6.x. To do so, perform the following:

1. Make a directory \TEMP on your hard disk.
2. Copy the file EXPAND.EXE from PC DOS 6.x diskette 1 into this directory, as follows:

```
COPY A:\EXPAND.EXE C:\TEMP
```

3. Copy these from PC DOS 6.x diskette 3:

```
COPY A:\CPBACKUP.* C:\TEMP
COPY A:\CPB*.OVL C:\TEMP
COPY A:\FILECHK.OVL C:\TEMP
```

4. Expand the following files:

```
CD \TEMP
EXPAND CPBACKUP.EX_ CPBACKUP.EXE
EXPAND CPBACKUP.HL_ CPBACKUP.HLP
EXPAND CPBACKUP.MS_ CPBACKUP.MSG
```

5. You can now run CPBACKUP from this directory to perform a backup of your system. The PC DOS 6.x CPBACKUP program can be used with PC DOS versions 4.0 and above, DR DOS 6.0, and MS-DOS 4.0 and above.

Remove Existing Compression, or

Coexist: The following sections explain how to use or remove any existing com-




Figure 6. DEFRAG Display Screen

pressed drives within your system in order to prepare for the installation of PC DOS 6.x SuperStor/DS compression. By its very nature, the information below may not be a complete list of all known compression systems. More information may be found in the PC DOS 6.x installation guide.

Stacker 3.0 or 3.1

This section describes how you can continue to use Stacker or migrate to PC DOS 6.x SuperStor/DS compression.

Stacker with PC DOS 6.x: You can upgrade a computer that has DOS Stacker already installed to PC DOS 6.x. Again, note that with the Stacker-compressed drives, you cannot use the SuperStor/DS compression drivers or utilities that are supplied with PC DOS 6.x.

To install PC DOS 6.x, make sure that the original DOS is running with the Stacker drivers loaded. (The original DOS must be one of the versions listed above in step 5 of the section "Back Up Your System.") Then, with PC DOS 6.x diskette 1 in drive A, enter A:\SETUP at a command prompt.

Remember *not* to select SuperStor/DS compression as an installation option, but select the other options you need and continue the installation.

Once the installation is complete, your computer is restarted, changes to the CONFIG.SYS file (or other files) are noted by Stacker, and you are prompted to reboot. All of this is quite normal! Stacker also then updates down-level files between the host drive and the compressed drive.

Stacker Removal: Before you begin the task of uncompressing your stacked (compressed) drives, make sure that you have an up-to-date backup of your complete system. The method described in the section "Back Up Your System" will help. It is also worth performing a CHKDSK on your hard disks to ensure that you will not encounter additional problems.

Unstack the Compressed Drives:

There are two ways in which you may uncompress the stacked drives, either from the command line or via the STAC utility.


```

Stacker Disk Mount Utility - 3.00, (c) 1990-92 Stac Electronics, Carlsbad, CA
Registered to:
    XXX XXXXXXXX
    ITSO

Stacker drive map:
  Drive A: was drive A: at boot time
  Drive B: was drive B: at boot time
  Drive C: was drive D: at boot time  [ D:\STACKVOL.DSK = 55.7 MB ]
  Drive D: was drive C: at boot time
Total of 1 available replaceable drive(s).

```

Figure 7. Screen Listing After Running STACKER Command

To use the STAC utility, enter STAC from the command line, and select the option UNSTACK.

To start the UNSTACK utility from the command line, enter UNSTACK X: where X is the drive you want to uncompress. If you are unsure which drive is compressed, enter STACKER from the command line. You will see the information in Figure 7.

In Figure 7, we see that we need to uncompress drive C, which is done by typing UNSTACK C: . Your particular setup may have multiple drives that also need to be UNSTACKED.

Whichever method you chose, at this point the two methods follow the same procedures.

The UNSTACK utility program now determines how much space is needed to uncompress the compressed files. It will let you continue only if there is sufficient space. You should check for any unwanted files that can be deleted, or you may need to back up and remove data until you are able to continue.

Important! The UNSTACK utility may take some time to complete, so it is imperative that you do not switch off your computer while the utility is progressing. If you do, you will lose some data.

Remove Old Files: You should now remove all files within the Stacker direc-

tory, then the Stacker directory itself. Also remove the REMed lines within CONFIG.SYS that refer back to the Stacker device drivers.

You are now ready to install PC DOS 6.x compression. See the section "Installation of PC DOS 6.x Compression" for more information.

Microsoft DOS 6.0 DoubleSpace

This section describes how you can use or remove the already compressed drives compressed by DoubleSpace. The following information pertains only to workstations that have the MS-DOS 6.0 upgrade package; we were not able to test workstations that were preloaded with MS-DOS 6.0.

DoubleSpace Coexistence: If you plan to use PC DOS 6.x SuperStor/DS on a computer that already has Microsoft DoubleSpace 6.0-compressed drives, you will not have to remove the DoubleSpace compressed disks first – PC DOS 6.x is able to use existing DoubleSpace-compressed disks.

However, note that there are a couple of hidden files that you should remove from your root directory after you finish installing PC DOS 6.x with SuperStor/DS compression.

1. Install PC DOS 6.x using one of these methods:
 - Boot from PC DOS 6.x diskette 1.

- Boot from hard disk; then, insert PC DOS 6.x diskette 1 into drive A, and from a command prompt, enter A:\SETUP .

Once the installation is complete, continue:

2. At the C: prompt, enter:

```

ATTRIB C:\IO.SYS -r -h -s
ATTRIB C:\MSDOS.SYS -r -h -s
DEL C:\IO.SYS
DEL C:\MSDOS.SYS

```

3. Remove the following line from the CONFIG.SYS file:

```

DEVICEHIGH=C:\DOS\DBLSPACE.SYS
/MOVE

```

This step is not essential, because PC DOS 6.x ignores this line, but it makes things look neater!

4. Check that all copies of COMMAND.COM on any of the drives are for PC DOS 6.x.

DoubleSpace Removal: There are no utilities supplied with MS-DOS 6.0 to uncompress and remove the DoubleSpace drives. However, once you have installed PC DOS 6.x, you can use the SSUNCOMP.EXE utility program. This utility attempts to uncompress the DoubleSpace drives, if there is sufficient space to do so. Ensure that the disk is free of any unwanted files, and that you can run CHKDSK cleanly. Again, before undertaking this procedure, it is worth backing up your system as described in the section "Back Up Your System."

Important! The SSUNCOMP utility may take some time to complete, so it is important that you do not switch off your computer while the utility is progressing. If you do, you will lose data.

DR DOS 6.0 SuperStor

The version of SuperStor** that is part of DR DOS 6.0 is not compatible with SuperStor/DS in PC DOS 6.x.

DR DOS 6.0 SuperStor Compression

Coexistence: This method is not recommended, because you will have down-level SuperStor drivers, but you may wish to use the existing DR DOS 6.0 SuperStor compression for a short time. Using it with PC DOS 6.x requires some manual intervention after installation of PC DOS 6.x.

We cannot stress too highly that you should *always* back up your system before attempting any major upgrade or system alteration.

To install PC DOS 6.x, make sure that the DR DOS 6.0 operating system is running and that the SuperStor drivers are loaded.

DR DOS 6.0 with SuperStor reserves a small area of disk that is uncompressed. When DR DOS starts, it treats this area as the C drive or boot drive. In this small area are the files necessary to boot the computer and the device drivers for compression. There must be an area of decompressed space in order for the following steps to work correctly.

1. Start the DR DOS operating system.
2. With PC DOS 6.x diskette 1 in drive A, enter A:\SETUP at a DR DOS 6.0 command-line prompt.

Remember *not* to select SuperStor/DS compression as an installation option (because we are trying to keep DR-DOS compression), but select the other options you require and continue the installation.

```

DEVICE=C:\SSTORDRV.SYS
rem Normal CONFIG.SYS file contents below
DEVICE=D:\DRDOS\HIMEM.SYS
DEVICE=D:\EMM386.EXE RAM
DOS=HIGH
.....
.....
rem Now Swap Drive C with D
DEVICE=C:\DEVSWAP.COM

```

Figure 8. Sample CONFIG.SYS File After Modification

Once the installation is complete and you are required to reboot, in all probability the system will not start.

3. Use PC DOS 6.x diskette 1 to boot from. When you are given the option of installing, cancel the installation, in order to return to a DOS prompt.
4. At the A:> DOS prompt, enter SYS C: .
5. Select drive C. At this point, you should see a small C drive. During normal DR DOS bootup, drive C is swapped to drive D, and drive D becomes drive C, which is the compressed drive.
6. Copy the file called DCONFIG.SYS to CONFIG.SYS. This DCONFIG file is the DR DOS configuration file that normally loads the compression device driver, swaps the boot drive for the compressed drive, then finally chains the configuration file on the compressed drive. This use of two configuration files will not work under PC DOS 6.x.
7. Remove the following line from CONFIG.SYS:
CHAIN=C:\CONFIG.SYS
8. Add any existing configuration lines from your original CONFIG.SYS file before the reference to DEVSWAP.COM. The DEVSWAP program is used to swap the drive letters of the boot drive and the compressed drive. If you have added more configuration lines, it is there-

fore important that those additional lines must point to valid programs on the D drive.

9. Save the configuration file, and reboot the system.

Figure 8 lists a sample CONFIG.SYS file after it is modified.

DR DOS 6.0 SuperStor Removal:

The version of SuperStor that is supplied with DR DOS 6.0 does not provide a way to decompress or uninstall disks once they have been compressed. In this case, before installing PC DOS 6.x, you will have to back up your system and delete or remove the drives.

1. Back up your system with a program that is compatible with PC DOS 6.x. See the earlier section "Back Up Your System."
2. At the DOS command prompt, enter SSTOR.
3. Select REMOVE from the menu, and follow the prompts. *Warning!* The REMOVE procedure of SSTOR deletes the compressed drive(s) and any data that they may contain; it does not uncompress or uninstall first.
4. To remove any reference to SuperStor in CONFIG.SYS, remove or REM these lines:

```

DEVICE=\SSTORDRV.SYS
DEVICE=\DEVSWAP.COM

```

You are now ready to install PC DOS 6.x with compression.

Installation of PC DOS 6.x Compression

Remember that in some countries, the compression software was shipped at a later time as an upgrade.

Complete Package: The installation of the SuperStor/DS compression utilities is straightforward – they are selected at the start of installation, along with the other tools and utilities. During the installation process, all the required programs and files are copied onto your target drive. At this point, you do not have a compressed drive.

If you wish to install SuperStor/DS at a later date, you may do so by inserting diskette 1 into drive A and entering A:\SETUP /E . This allows you to select and install the PC DOS 6.x tools only, without having to reinstall the entire operating system again.

Separately Shipped Compression

Upgrade: As part of the PC DOS 6.x compression upgrade, you receive two diskettes – a replacement diskette 1 and a new diskette 5 (the compression diskette).

If you have already installed PC DOS 6.x without compression, it is a very simple matter to add the compression tools.

1. Start the current PC DOS 6.x from the hard disk.
2. Select the DOS command prompt.
3. Insert the new diskette 1 into drive A, and enter A:\SETUP /E . The /E option selects just the tools to be installed, and not the entire base of DOS files and utilities.

At this point, the SuperStor/DS compression driver has not been enabled. To enable it, you must use the SSTOR.EXE utility. The first invocation of SSTOR (with the option PREPARE) detects that compression is not enabled, and will enable it for you. Your computer then reboots.

SSTOR creates two files, one the hidden file DBLSPACE.INI, and the other file DBLSPACE.BIN in the root of the boot drive.

Compressing Your Hard Disk or Diskettes

Or, how to put a quart into a pint container!

Before compressing your hard disk, you should take time to create a bootable DOS diskette. This diskette can be used to recover from hard-disk problems in the unlikely event that they occur.

Recovery Diskette: This is a simple, but very important, procedure.

1. Insert a blank 720 KB or larger diskette into drive A.
2. Enter FORMAT A: /S /U to create a bootable system diskette.
3. From the A: prompt, copy the following files onto this diskette:

```
COPY C:\DOS\E.EXE
COPY C:\DOS\E.EX
COPY C:\DOS\ATTRIB.EXE
COPY C:\DOS\CHKDSK.COM
COPY C:\DOS\RTTOOL.EXE
COPY C:\DOS\DEFRAG.EXE
COPY C:\DOS\FORMAT.COM
COPY C:\DOS\SSTOR.EXE
COPY C:\DOS\SSUTIL.EXE
COPY C:\DOS\DBLSPACE.BIN
```

4. On the bootable system diskette, create a CONFIG.SYS file that contains only the line FILES=30 .
5. Try booting from this diskette to ensure that it works.

This diskette contains all the utilities you need if a problem occurs with your compressed hard disk.

SSTOR Utility

Compressing your hard disk could never be more simple! All you need to do is run the utility SSTOR.EXE.

The SSTOR utility program is menu- or command-line-driven, and has the following options:

Preview, which reads the data on your entire disk and reports how much it can be compressed. Preview informs you what compression ratio you may achieve; you can then use this number during the PREPARE option.

Prepare, which creates compressed disks. You may either compress an entire disk along with the files on it, or make a new compressed disk from the free space on a drive.

Mount, which chooses a compressed drive to mount. When a compressed drive is mounted, you can access the data in that drive by referencing it with its drive ID. Mounting a compressed drive “links” it to a drive ID. The compressed drive must have been created by the PREPARE option of SSTOR.

Unmount, which unmounts a SuperStor/DS-compressed drive. Once the drive is unmounted, you cannot access the data on it until you mount the disk again.

Statistics, which displays the disk and compression statistics for the currently active compressed drive(s).

Password. You can install a password on a compressed disk to prevent unauthorized use. *Warning:* You must remember your password in order to access your disk. There is no way to recover a lost password, so it is imperative to save it on a diskette when you are prompted to do so, and to put away the diskette for safekeeping.

Delete, which removes a compressed disk. *Warning:* DELETE results in the loss of all data on that disk. If you DELETE a compressed hard disk, all data on it is destroyed. If you DELETE a mounted disk or automounted diskette, the disk will be unmounted and all data deleted. You cannot recover deleted data with the DOS UNDELETE utility.

Compressing Your Diskettes

As well as increasing the capacity of your hard disks, using SuperStor/DS allows you to create diskettes that have increased capacity. (Note, however, that you cannot compress diskettes of capacity 360 KB or 720 KB.)

A unique feature of SuperStor/DS compression is its ability to mount automatically compressed diskettes. This capability offers the great advantage of not having to manually mount and unmount each diskette as it is placed into the diskette drive.

Portable Diskettes: You can create two types of compressed diskettes:

1. Diskettes used only on PC DOS 6.x systems with SuperStor/DS compression. To create such a diskette:
 - A. Place a FORMATTED diskette into your diskette drive.
 - B. Enter SSTOR at the DOS command prompt.
 - C. Select the PREPARE option.
 - D. Select "Compress floppy disk."

At this point, you should choose which diskette drive you are using, then the compression ratio you wish to use. In most cases, use the standard ratio of 2:1, unless you are transferring a quantity of similar file types.

 - E. Confirm that you do *not* want to make this diskette a Universal Data Exchange (UDE) diskette.

Your diskette is now ready for use.

2. Diskettes used to transport data to computers that have other levels of DOS, or do not have SuperStor/DS compression. To create this kind of diskette, the procedure is basically the same as the one above, except that in step E you should choose to make this diskette a UDE diskette. The UDE option puts a simple compression driver onto the target diskette.

Having been prepared with the UDE option, the diskette has some space left uncompressed; this space can only be seen or used by DOS systems that do not have SuperStor/DS (or Microsoft DoubleSpace) compression loaded. This uncompressed space contains the utility UDEON.COM, which is a Terminate and Stay Resident (TSR) program. When run on a DOS system, this utility permits read/write access to the compressed part of the diskette. *Note:* UDEON.COM does not function under the Multiple Virtual DOS Machine (MVDM) support in OS/2 2.1 and earlier versions.

When the UDE utility is no longer required, it may be turned off by running the UDEOFF.COM program, which is found in the compressed part of the diskette.

Commands Used With Compressed Drives

Here is a list of PC DOS 6.x commands that pertain to compressed drives:

- **ASSIGN.** The ASSIGN command should not be used on a drive that has been compressed, or on the resulting compressed host drive. In fact, it returns an error if you attempt to use it.
- **CHKDSK (RTOOL).** CHKDSK has been enhanced to call the RTOOL program, which performs an integrity check on compressed drives. It is able to repair faults in a way similar to the repair of normal drives.
- **DIR.** The DIRectory command has been enhanced to include additional switches for showing compression information about the files and directories on the compressed drive(s).

The /C switch displays compression ratio for files and directories. This switch is ignored if it used in conjunction with the /W or the /B switch.

The /O switch lists files in sorted order. When used with the /C switch, the files are sorted based on compression ratios.




```

...
Volume COMPRESSED created 10-01-1993 4:42p
Volume Serial Number is 37ED-BCD5

Found lost cluster #1058, MDFAT should mark for deletion

MDFAT/FAT crosscheck found 1 lost cluster(s).

***** Disk Usage Summary *****
...

```

Figure 9. CHKDSK Output After Turning Power Off Prematurely

- **FORMAT.** The FORMAT command can be used on a compressed drive. All data on that compressed drive will be deleted. The compressed drive itself still exists, occupying the same space. Once formatted, it can store new data, in compressed form. *Note:* If you use FORMAT on a compressed diskette, it will also remove the UDEOFF.COM command from the diskette.
- **JOIN.** The JOIN command should not be used on a drive that has been compressed or on the resulting compressed host drive. In fact, it returns an error if you attempt to use it.
- **MOUNT.** The MOUNT command can be run from the DOS command line or via the SSTOR program as a menu item. MOUNT is used to assign a specific drive letter to a SuperStor/DS disk. The letter remains in effect until the drive is UNMOUNTed or the system is restarted. With its /V option, MOUNT can also be used to display mountable drives.
- **SSUNCOMP.** SSUNCOMP is a utility program that uncompresses drives which were compressed with either SuperStor/DS or DoubleSpace. This utility is menu-driven.
- **SSTOR.** SSTOR enables you to preview a disk, prepare a disk or diskette for compression, mount a compressed drive, unmount a drive, show statistics about a drive, or set a password for a compressed drive. This program is menu- and command-line driven.
- **SSUTIL.** The SSUTIL program runs a set of utilities designed to recover from disk problems, tune the disk by defragmenting, or alter the size of the drive.
- **UDEOFF.** This is a small utility that unloads the UDEON TSR.
- **UDEON.** This small TSR program allows the use of compressed diskettes on a DOS machine that does not have SuperStor/DS installed.
- **UNMOUNT.** The UNMOUNT command can be run from the DOS command line or via the SSTOR program as a menu item. UNMOUNT releases a mountable SuperStor/DS disk from its assigned drive letter. This drive is then not accessible until it is MOUNTed again. With its /V option, UNMOUNT can also be used to display mounted drives.

Recovering from Problems

From time to time, you may receive error messages or have problems with your hard disk. These things could be caused by a sudden power loss, by switching your system off while there is still disk activity, or even when an unrunly program terminates abruptly.

Hard Disk Does Not Boot: When your hard disk does not boot, you are suddenly grateful that you created a

bootable DOS diskette – the one you created before or after you made your first compressed drive. This diskette contains all the tools required to fix a disk problem. Some of these tools are:

- **CHKDSK.COM**, which reports and fixes cluster problems on your uncompressed drives.
- **RTOOL.EXE**, which is invoked by CHKDSK to report and fix cluster problems.
- **SSUTIL.EXE**, a set of compressed-drive utility programs that enable you to check, repair, and scan for disk problems. If problems are found, you will have choices for recovery.

CHKDSK Finds Errors: Occasionally you may run CHKDSK to determine the status of your hard disk. CHKDSK checks the uncompressed portions of your hard disk, and then it invokes the RTOOL program to check the compressed part of your disk.

As an example, we started a copy operation and then turned the power off halfway through. After restarting and rebooting, we ran CHKDSK, and saw the screen in Figure 9.

In cases such as this, you can run CHKDSK /F against this drive to fix the problem.

Superstor/DS Compression and OS/2

At the time this article was written, the PC DOS 6.x SuperStor/DS compression did not function within the DOS environments in OS/2 1.x, 2.0, or 2.1 – the DOS Box, a Multiple Virtual DOS Machine, or a Booted Virtual DOS Machine.

Also, the UDEON/UDEOFF utilities did not operate correctly in an OS/2 environment, so you cannot copy data on compressed diskettes from your PC DOS 6.x system to an OS/2 2.1 or earlier system.

OS/2 Workplace Software Strategy: Creating New Possibilities

Mary Forlenza and Cheryl Leone
IBM Personal Software Products
Communications
Boca Raton, Florida

IBM Workplace is not an operating system. Instead, it is a set of new operating system technologies. This article discusses the major entities in Workplace, entities that have the common theme of openness – open standards; vendor-neutral technologies; an open, object-oriented architecture; an enabling technology for sharing objects; and open, distributed computing. It points out how businesses can use these Workplace entities to respond quickly to industry changes. It also discusses how the IBM microkernel will make it possible to port applications across several hardware platforms. In essence, this article is about IBM's Personal Software Products strategy, the use of Workplace technologies that will bring about people-oriented computing.

The Software Solution That Meets Your Needs Today and Tomorrow

Making decisions about which hardware and software to buy is daunting when about 200 new products are announced every week. When shopping for products, we all look for software that is compatible with our current hardware and software. And we want to keep up with new technology that promises to provide better ease of use, more capabilities, and overall efficiency to our business systems. But we worry about making the right choices so we can continue to use the products we have today as we move toward future technology.

IBM is currently building upon the compatibility, usability, and advanced technology that Operating System/2 (OS/2)* offers. Customer feedback has helped IBM draw plans for future software that will enable you to:

- Access data through a more natural interface for greater productivity
- Share information across departments and locations
- Capture new business opportunities using timely software solutions
- Easily integrate, use, and manage networks
- Protect your investment in hardware, software, and skills
- Implement technology at your own pace as your business or personal needs grow

IBM Personal Software Products has formulated a strategy for developing products that will make use of the IBM Workplace technologies. Perhaps you have heard or read about IBM Workplace. This is our name, not for an operating system, but for a set of new operating system technologies that will fill your needs by making our systems more portable, interoperable, and scalable.

Workplace technologies are innovations that IBM will share across the OS/2, OS/400*, and AIX* products. IBM will integrate state-of-the-art elements that are necessary for truly seamless distributed computing. At the same time, IBM will support and enhance these products, protecting your investment in existing systems and applications.

This article describes how the Personal Software Products strategy will use IBM Workplace technologies to fill your needs with specific software solutions.

Workplace – Based on How You Use and Share Information: Workplace is designed around the way you work, whom you share information with, and

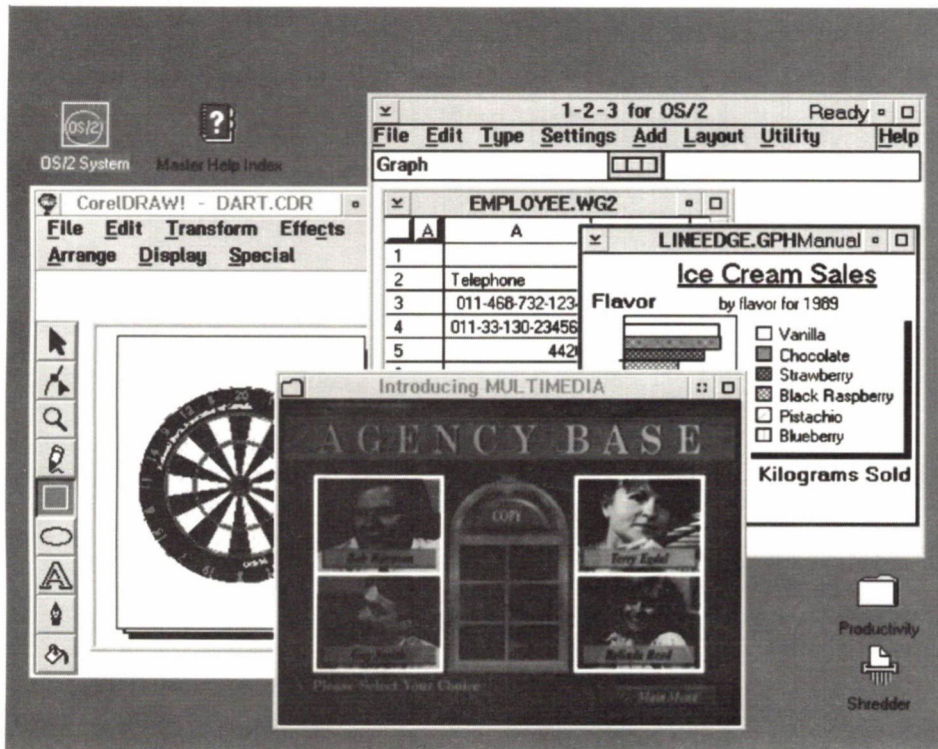
where your information is located. Its standards and technologies resolve much of the complexity caused by the diversity of hardware and software available today for personal computers. A major aim of Workplace is to simplify the decisions you must make when you buy hardware and software.

Workplace's Open Standards Give You Choices: IBM Personal Software Products will provide software that is usable on multiple hardware platforms, such as both Intel** and RISC microprocessors, and with applications from multiple vendors. Workplace can accomplish this because it combines IBM technology with open industry standards supported by key industry vendors. The important thing to remember about IBM Workplace is that it is *not* proprietary. You will benefit by being free to choose the best products in the marketplace for your particular business or personal needs. IBM Workplace will enable you to combine products to form the best solutions for your needs in ways not possible before.

Elements of the Workplace Solution: The Workplace solution offers products that contain these elements:

- Consistent and natural user interfaces that enhance your productivity
- High-quality, reusable software delivered to users more quickly through object technology
- Local area network software that facilitates sharing information with co-workers and is easy to manage across multivendor hardware and software
- Software that is portable across multiple hardware platforms to accommodate the most diverse business systems environment

The following sections explain each of these elements, including what is available now and the plans for the future.



An Interface You Don't Have to Think About

OS/2 contains an advanced object-oriented user interface called the *Workplace Shell*®. This interface is "so obviously useful that you stop thinking about it at all," wrote James Fallows, an editor for *The Atlantic Monthly* (February 1993). The OS/2 Workplace Shell currently provides a single interface to DOS, Windows, and OS/2 applications.

A common interface, based on the OS/2 Workplace Shell, will be adopted across Workplace products, including DOS and UNIX®. This consistent interface will enable you to have a single, familiar way to view and manipulate data from different applications, operating systems, and hardware types.

The Workplace Shell graphical user interface is object-oriented to be people-oriented; that is, it reflects the way a person works in the real world. You can just work with familiar objects, such as

documents, file folders, and shredders, rather than having to learn commands and deal with programs. By working directly with the data your business requires, rather than concerning yourself with the programs using that data, you can truly use your personal computer as a productivity tool.

The Workplace Shell already provides true object manipulation, while the competition is still trying to develop it. Its icons represent objects you can work with directly; for example, to print a message, just drag the object or icon of a Lotus cc:Mail® in-box message to the OS/2 desktop printer. The Workplace Shell even provides object management for an entire project, which means that if you open the folder for a work project, every file and application you were last working with will reopen – at your last stopping point. Peter Coffee wrote in *PC Week* (June 14, 1993) that OS/2 2.1 "performs operations with an elegant user interface that has fewer nooks and crannies than Windows, but delivers

greater control over system behavior than the Macintosh®."

Paul Somerson, editor-in-chief of *PC/Computing*, said at COMDEX® in November 1993, "OS/2 has the most attractive user interface available on the PC . . ." (*PC/Computing* named OS/2 the Most Valuable Product in the Systems Software category for the second consecutive year at COMDEX. OS/2 has won more than 35 major industry awards worldwide.)

Advancing from Computer-Oriented People to People-Oriented Computing:

IBM Workplace will provide more natural user interfaces to the computer, so that you could, for example, write a message with a pen, and use a voice command to send the message to a co-worker. By extending operating systems with multimedia technology, speech recognition, and pen and touch technology, Workplace will create a natural computing environment where you can more fully use your human abilities. Natural computing accommodates new users who may not be familiar with a mouse or keyboard, or cannot use them under certain circumstances, such as with medical applications. As the computer becomes more adapted to natural human interaction, it will become a more useful tool.

Human-centered technologies have begun to appear in IBM products. These technologies will be employed in Workplace products in the future. The IBM Person to Person® product, available today, lets as many as eight people collaborate on a project simultaneously by viewing the same document file and annotating their comments with a pen or keystrokes, or both. The pen is used to execute commands by touching areas of the screen. Another IBM product, the Personal Dictation System, has a vocabulary of 32,000 words and recognizes words as fast as a good typist can type. With the Personal Dictation System you can command and control your OS/2 desktop with your voice; for example,

you can select icons, open folders, and move windows. This product was a *Byte Magazine* Best of Show finalist for the Best New Software at COMDEX in November 1993.

When you can get work done by merely talking to your computer, you will spend less time in training or reading instruction manuals. You can concentrate on getting your work done rather than on managing your computer.

OpenDoc Lets You Work with Documents, Not Applications: In the future, Workplace will include a compound document architecture that allows for a mixture of many data types within a single document. You will be able to manipulate data in *documents* rather than in different applications. This object-oriented architecture, called *OpenDoc***, will enable you to create powerful multimedia documents that mix data types such as text, graphics, audio notation, and video clips within a single document.

Today's monolithic applications grow more difficult to use and maintain as functions are added. In the future, smaller *parts* will replace traditional application functionality. For example, the standard components of a word processing package would be a combination of text parts, spreadsheet parts, and a spell-checking part. *Part handlers* will work behind the scenes to enable the parts to be used in a single document and appear seamless to the user. OpenDoc will take care of the complexity of combining different data types and shield the user from having to manipulate diverse application programs.

OpenDoc will enable you to view and edit data types, such as charts and spreadsheets, at the same time in a single document window rather than switching between application windows. This is called *editing in place*. You will be able to manipulate them using cut-and-paste or drag-and-drop capabilities.

Industry Supports OpenDoc

OpenDoc is supported by key vendors in a non-profit company called Component Integration Laboratories. The founding companies of CI Labs are:

- Apple
- IBM
- Novell
- SunSoft
- Taligent
- WordPerfect
- XSoft

OpenDoc will be available on multiple platforms: initially OS/2, Macintosh, UNIX, and DOS/Windows. OpenDoc will interoperate with Microsoft's Object Linking and Embedding (OLE) 2.0 specification and, over time, with the Taligent compound document model.

Today you might have one way to create charts in your spreadsheet application and a different way to create charts in your presentation program. In the future, with OpenDoc, you will select the chart program you prefer, and use it for the charts in all of your documents.

OpenDoc will provide a standard way to share document parts between programming languages and across operating systems and networks.

With Objects, Businesses Can Respond Quickly to Industry Changes

The impact of objects will be far-reaching. While an object-oriented user interface is one in which the user works with familiar objects, object-orientation also refers to the objects or modules of code a programmer uses to create a product.

After being on the sidelines for more than 25 years, object technology has entered the mainstream of computing because of business needs and rapid technological innovations. The effect on software is equivalent to the radical change in the hardware industry that was triggered by PCs, which allowed vendors to build systems from standard "off-the-shelf" parts.

IBM and other companies are adopting object technology to greatly reduce the time required to develop system and application software. This, in turn, will enable businesses to respond quickly to industry changes and become more competitive.

Objects are Reusable Software

Components: The concept behind *object technology* is to solve a complex business problem by assembling and/or extending reusable software components, or objects. An *object* is defined as the data and logic that represents a useful element in an application. For example, in a financial application, there may be objects that represent account, branch, and customer. The account object can be defined to perform certain functions, such as open, deposit, or withdrawal; and, the programmer doesn't need to understand the complex details of how the object works. The value of objects in designing and implementing software is that they hide complexity, enable the re-use of standard components, and simplify the programmer's task.

Object Frameworks Perform Specific

Tasks: Programmers will be able to reuse existing objects and sets of objects to build programs far more quickly and easily. Sets of objects that perform specific tasks are called *frameworks*. For example, graphical user interface frameworks let objects like windows and menus and dialog boxes work together to provide an efficient way to build consistent graphical applications. Another simple example would be adding fax and spreadsheet capabilities to a financial application program by customizing

Objects Begin with SOM

IBM's object technology begins with the System Object Model* (SOM), available in OS/2 and AIX today. SOM defines an enabling technology for sharing objects. It allows objects to be packaged in a way that exposes only their interface. As a result, an object can be written in one language, such as C++, and used or refined by another language, such as REXX. With SOM in OS/2, developers are able, for the first time in the software industry, to use objects written in one language with programs written in different languages without requiring the object's source code. Developers can build object-oriented applications today by using the SOMobjects Developers Toolkit for OS/2, AIX, and DOS/Windows, or by using any of the many SOM-enabled development toolkits from other industry leaders.

SOM supports transparent remote access to objects in a distributed environment. SOM is fully compliant with the open, distributed objects specification adopted by the Object Management Group (OMG),** an industry consortium. This specification also has been adopted by the Open Software Foundation (OSF),** and has substantial backing from the industry, including these major vendors:

Hewlett-Packard	Sun	Apple
Borland	Lotus	Intel
WATCOM	Novell	Univel
AT&T/NCR	SunSoft	Metaware
DEC	UNIX System Laboratories	

IBM is implementing SOM across its operating systems; thus, OS/2, OS/400, AIX, and MVS will have the SOM enabling technology. Other leading vendors also are implementing SOM in their systems, including all of the OpenDoc founding companies listed previously.

SOMobjects supports all of the important industry network transports: TCP/IP, Novell's IPX, and NetBIOS. It will take advantage of OSF's Distributed Computing Environment (DCE) network services to provide location transparency, enhanced security, and support across multivendor hardware and software.

existing fax and spreadsheet frameworks. There are frameworks for specific business problems, too, such as hospital management or inventory control. Well-designed frameworks provide the general design and implementation for a specific problem and allow a developer to customize them to fit a particular situation.

Solving Business Problems: Here are some other benefits of object technology:

- Applications more closely model the business problem they are designed to solve, thus they are easier to develop and maintain. Conversations center around the loan approval or portfolio, rather than the XYZ transaction, as it might have been called in a computer-centered implementation.

- The data and actions to be taken on that data are bound together and self-contained; thus, the personal banker doesn't need to know the computer steps or logic, but just the loan approval process. If a change had to be made to accommodate a new law, for example, then the loan approval object could be customized, enabling the banker to automatically take advantage of the update. The banker's productivity would be enhanced.
- Since frameworks are flexible and easily extended, programmer effort is leveraged over and over again. Fewer lines of programming code need to be written to build applications. This means far less time is spent developing and testing programs. Existing frameworks have already been tested for quality. And programs can be delivered to the user faster, thus improving companies' time-to-market and competitive posture.

Freed from having to program complex components of the operating environment, developers can solve business problems more innovatively and at lower cost. They will spend more time on their own area of expertise, such as financial or manufacturing functions. Applications will be easier to maintain. This, in turn, will provide users with higher-quality, faster solutions to problems. Studies have suggested 10- to 100-times productivity gains in writing and maintaining programs using object-oriented technology.

Frameworks used today in the SOMobjects* Developers Toolkit for OS/2 and AIX will be compatible with frameworks available in the future. For re-use of software components to be effective, the goal must be to enable re-use across platforms, vendors, and tools.

Taligent Will Provide Object

Frameworks: Taligent, Inc., a joint venture between IBM, Apple**, and Hewlett-Packard**, will be the source of many future object frameworks. Businesses

and software developers will use these object frameworks to create leading-edge software solutions. Application developers can make use of the Taligent** application frameworks, and know that their application programs will run on a variety of operating systems.

Advancing to Open, Distributed Computing

IBM has made distributed computing one of the company's biggest investments. "We're going to help our customers solve the problems of distributed computing – by making connectivity open and easy across IBM platforms and between IBM and competitors' platforms," said IBM CEO Lou Gerstner in a March 1994 speech.

To accomplish this goal, IBM recently defined a set of guidelines for deploying open, distributed systems. IBM's guide to distributed computing is called *Open Blueprint*. Our strategy is to advance customers and technology from today's file and resource sharing to true distributed computing. This will enable you to distribute data and functions where they are most needed and make full use of the processing power across a network.

Workplace offers a comprehensive local area network (LAN) product line today to connect users, whether they are in small workgroups or throughout large enterprises. Our strategy is to provide *alix scalable* products, meaning that as your networking needs grow, you can add on products, without having to replace the software that you already have. You can evolve gradually to distributed computing by "snapping on" components one by one. As you enlarge your network, you will also gain increased functionality, such as enhanced systems management and security.

LAN Manages Complexity So You Can Manage Information: IBM LAN systems currently provides a manageable environment and gives the appearance of a single system. You can access

the information you need without dealing with the complexities of the network. As your network adds more users, and multivendor hardware and software, Workplace products will provide all of the function necessary for you to continue to work with network resources as though they were on your personal desktop system.

IBM will support this distributed environment with technologies from standards groups, such as OSF. To protect your investment in hardware and software, IBM will maintain its lead in the industry for compliance with open, vendor-neutral, industry standards. IBM will also enhance its own distributed products that have earned broad customer acceptance, such as our DB2* family of relational databases, our Communication Manager/2, and our CICS* family of transaction processing systems.

Many of IBM's major software products that run on our key operating systems already conform to standards described in the Open Blueprint. Several products that will provide the path to distributed computing are described on the next page.

LAN Systems Offers Workgroup

Applications: Electronic mail and workgroup computing are major elements of any networked PC environment. IBM gives you a simple but powerful entry into workgroup computing with the IBM Advanced Server for Workgroups. This suite of groupware functions, centered around IBM LAN Server and Lotus Notes**, provides knowledge-sharing applications, mail and messaging capabilities, a development environment for distributed applications, and discussion and collaboration databases.

The IBM Advanced Server for Workgroups is ideal both as an entry point for workgroup computing over the LAN and for large corporations needing to share applications over interconnected networks around the world.

Single-System Image with DCE:

IBM's DCE products for OS/2, AIX, and DOS/Windows already support OSF's Distributed Computing Environment remote procedure call to invoke remote applications as if they resided on the same physical machine. IBM will integrate DCE directory and security services with LAN Server to provide a single-system image to the users on a network of multivendor computers and operating systems.

IBM is implementing DCE services across its operating systems; thus, DCE products will also be available for OS/400 and MVS.

Using the LAN as a System

Here is an example of one customer that benefited from IBM distributed computing technology. University of North Carolina Hospitals had their patient information residing on many different architectures that were added to accommodate the needs of different functions, such as the nurse's station, operating rooms, and laboratory. They asked IBM to devise a way to eliminate the need to log on to six different systems to get information about a single patient. The servers involved were from IBM, DEC**, Data General, and Stratus. We worked with them to put together an interoperable environment they could access with a single logon. This reduced the time it takes to get the information on any patient from two hours to a matter of minutes.

By moving to IBM LAN systems solutions today, you can begin a gradual shift to distributed computing. You can add products at your own pace, while continuing to use your existing software and hardware. The benefit will be your ability to distribute data and function

throughout your business, making them available to *anyone, any time, anywhere*. You will also be able to interconnect your systems with those of your customers, suppliers, and business partners. IBM's vision of a distributed computing environment with a single-system image is called "LAN as a system." We are applying this concept to Workplace products today to make your distributed computing environment more scalable, manageable, secure, and powerful.

LAN Systems Products:

LAN Server + OS/2 = The Premier Client/Server Offering: IBM LAN Server manages the resources that are distributed across the network. LAN Server extends OS/2's management of local resources to network resources. LAN Server is an application that exploits the power of OS/2 on server systems.

A *server* is a network system that receives requests for sharing files, devices, and applications from a user's system, which is called a *client*. OS/2 with LAN Server has become the premier client/server combination in the industry by giving users access to network resources, and enabling users to view them, with Workplace Shell, as though all resources were on their local workstation. In the future, LAN Server will become an even stronger base for distributed computing by fully utilizing Open Blueprint guidelines.

LAN Distance – Your Node on the Road: IBM LAN Distance* is available today to support mobile users who need to connect into their office LAN from remote sites anywhere in the world. LAN Distance enables you to access your LAN over public telephone lines and to run LAN-based applications. It supports all of your current client software, running remotely, with the same look and feel as your regular desktop.

NetView Manages the Works: IBM LAN NetView* is a systems manage-

ment family of products available today. The LAN NetView family enables you to monitor and control information processing resources, both hardware and software, either on a single LAN or throughout an enterprise having hundreds of users on multiple networks. LAN NetView also enables you to manage productively – either attended or unattended – software configuration, installation, and distribution (CID).

LAN NetView provides network administrators with a single, integrated view of LAN resources. Products in the NetView family monitor desktop and server performance, manage hardware and software inventory, and handle problem management in a distributed network environment.

Reach Multiple Platforms with the IBM Microkernel

Compatibility reached a higher level when IBM introduced OS/2 in March 1992. OS/2 is an operating system for Intel microprocessor-based PCs that exploits the high performance of today's 32-bit PC hardware. More than 1,300 IBM-compatible PC models are currently supported. OS/2 runs DOS and Windows applications as well as more than 2,000 native OS/2 applications.

Today's compatibility of hardware and software will evolve to portability tomorrow with Workplace technologies. The IBM microkernel-based architecture will enable you to run your application programs on different hardware processors. These could be Intel- and RISC-based systems you already have, or systems you want to buy, such as the PowerPC*, that will improve the speed and performance of your current software. Applications written to different operating systems will be able to run together on the same machine, permitting the use of exactly the right set of programs for a particular business or organization.

The microkernel-based architecture also solves the problem of having to rewrite

operating systems for different hardware processors. In the future, the new IBM microkernel will enable OS/2 to become available on RISC systems, while also enabling OS/2 to continue as our premier 32-bit Intel offering.

How the Microkernel Makes Software Portable

The IBM microkernel-based technology enables the operating system to be independent of the type of hardware processor. This is accomplished by rearranging the functions of today's traditional operating system into simplified, modular building blocks. The most essential functions of an operating system are contained in the "microkernel" module, and only a small part of that module is hardware-dependent. Other operating system services are contained in separate modules that are not hardware system-dependent. Examples of these services are database engines, file systems, and device drivers. These pieces can be developed by different vendors so that users can plug in their favorite database engine, for example. Operating system building blocks interact with the microkernel and services modules; and, unlike current systems, they also are independent from the hardware system.

The microkernel-based structure has another advantage – it is scalable from the supercomputer to the hand-held personal digital assistant, and even to the computer in your future car.

Operating Systems Become Personalities

This new architecture will use a microkernel "core" product to run multiple, diverse operating systems as "personalities" on a variety of hardware platforms. The *personalities* contain the elements that users recognize as a particular operating system today, including the look and feel of the system and the ability to run a specific set of applications. The initial personality will be OS/2 – with support for DOS and Windows. This personality, and others added in the future, will enable users to run a broad

range of existing and new applications on the same machine. Businesses will benefit by protecting their investment in hardware, software, and skills. Operating system services, called *common personality services*, will be packaged with personalities to accommodate specific markets. The services can be common and shared among personalities.

The combination of the IBM microkernel and the OS/2 personality will deliver OS/2 on other systems, such as the PowerPC. Developers will be able to exploit RISC technology via the OS/2 32-bit application program interface, which is the same API used with OS/2 on Intel-based PCs today.

A Bonus for Users – Software

Developed Faster: The IBM microkernel-based architecture will reduce the


IBM Microkernel Graduates from College

The IBM microkernel is derived from the Mach research project, which was conducted at Carnegie Mellon University, and from other research. The IBM microkernel takes advantage of the best recent research work, while providing an industrial product for systems vendors to build upon. It consists of a small privileged kernel plus a set of personality-neutral services for writing device drivers and other services, such as database engines, that can be used by a number of operating system environments.

time it takes system developers to create new software functions or to make use of new hardware functions. This will provide users with new technology more rapidly than in the past. In addition, the microkernel-based architecture breaks the traditional operating system into modular pieces that lend themselves to object-oriented implementation techniques.

Be Free To Choose the Best Products for Your Needs

In summary, IBM Personal Software Products is basing its strategy for delivering software products on Workplace technologies. Key elements of Workplace – the common user interface, object technology, and distributed computing services – are already available in OS/2 and LAN systems products.



Elements	Benefits	Technologies and Products
A consistent and natural user interface	For usability and productivity	OS/2 Workplace Shell with integrated multimedia, pen, voice, and touch
Object-oriented programming interfaces	For reusable, high-quality programs developed faster with less cost	SOMobjects Toolkit, OpenDoc, Taligent frameworks
Local area network software	To connect small workgroups through large enterprises using both non-IBM and IBM hardware and software	LAN Server, LAN NetView family, LAN Distance, Advanced Server for Workgroups, DCE for OS/2
Portable software	To run the same software on both Intel- and RISC-based hardware	IBM microkernel plus portable operating systems, including OS/2, DOS, and DOS/Windows

Figure 1. Workplace Elements, Benefits, Technologies, and Products

An Interface You Don't Have to

Think About: The Workplace Shell provides a natural user interface that is consistent across dissimilar applications. The future interface will become more people-oriented through expansion of multimedia technology, speech recognition, and pen and touch technology.

With Objects, Businesses Can Respond Quickly to Industry Changes:

Object technology has the potential to spark a new cycle of software development creativity as profound as the move to high-level programming languages more than 30 years ago. Some advantages of object technology are improved productivity for programmers and users, and higher quality software. Other benefits are shortened development cycles that keep pace with hardware advancements and enable you to respond quickly to changes in your industry.

Advancing to Open, Distributed

Computing: IBM has superior connectivity because our years of experience with networking enables us to support a broad range of products and offer our own comprehensive product line for client/server computing. We are investing heavily in distributed computing to make communications across networks of PCs more powerful, manageable, and secure for implementers, administrators, and users.

Reach Multiple Platforms with the

IBM Microkernel: The IBM microkernel-based architecture will greatly simplify systems software and enable it to be ported across a variety of hardware platforms. The goal is to enable you to select the best software and hardware for your needs from multiple vendors, and have one interface to your applications and even to a network of systems.

The use of open industry standards and vendor-neutral technologies in the Workplace strategy ensures that the hardware and software you buy, and the applications you build today, will be usable with Workplace products in the future.

You can get started right away, but you don't need to add new technology all at once. You can progress at a pace that meets the requirements of your growing business or your personal computing needs.

The elements and benefits of Workplace, and relevant technologies and products, are listed in Figure 1.

Getting Started Today with Workplace

Users can start to take advantage of the IBM Workplace standards and technologies used by Personal Software Products

by installing OS/2 and OS/2 LAN Server today. To order these products, call OS/2 and LAN Systems Sales at 1-800-3-IBM-OS2 in the U.S. and 1-800-465-7999 in Canada. For technical support on OS/2 and LAN software products, call the IBM Technical Support Center at 1-800-992-4777.

Developers wanting to obtain the tools they need to take advantage of an object-oriented environment can call The Developer Connection for OS/2 at 1-800-6DEVCON.

Mary Forlenza is a writer and editor in IBM Personal Software Products Communications, Boca Raton, Florida. Mary has a BS degree from the University of Florida and is completing an MA degree in Communications at Barry University. She can be reached via Internet at maryf@bcvml1.vnet.ibm.com.

Cheryl Leone is a program manager in the IBM Personal Software Products marketing organization, Boca Raton, Florida. She is responsible for developing marketing strategies for IBM PSP products. Cheryl has a BA degree from Furman University.

Trademarks

The following terms, denoted by an asterisk (*) in this publication, are trademarks of the International Business Machines Corporation:

IBM, AIX, CICS, DB2, LAN Distance, Micro Channel, NetView, Operating System/2, OS/2, OS/400, Person to Person, PS/2, PowerPC, Presentation Manager, System Object Model, SOMobjects, Superstor/DS, and Workplace Shell.

The following terms, denoted by a double asterisk (**), are trademarks of the companies noted:

Apple and Macintosh – Apple Computer, Inc.
Borland – Borland International, Inc.
COMDEX – Interface Group, Inc.
CompuServe – CompuServe, Inc.
DEC – Digital Equipment Corp.
Hewlett-Packard – Hewlett-Packard Company
Intel – Intel Corporation
Lotus, Lotus Notes and cc:Mail – Lotus Development Corp.
Metaware – Metaware, Inc.
Microsoft and Windows – Microsoft Corporation
Novell – Novell, Inc.
Object Management Group and OMG – Object Management Group, Inc.
Open Software Foundation and OSF – Open Software Foundation
PCMCIA – Personal Computer Memory Card International Association
PKZIP and PKLITE – PKWARE, Inc.
RAMBoost and Central Point Backup – Central Point Software, Inc.
Scrabble – Spears Games
Stacker – Stac Electronics
Sun and SunSoft – Sun Microsystems, Inc.
Superstor – Adstor, Inc.
Taligent – Taligent, Inc.
UNIX – X/Open Company Ltd.
WATCOM – WATCOM
WordPerfect – WordPerfect Corporation
XSoft – Xerox Corp.

All other product names may be trademarks of their respective companies.

“PC DOS 6.3's memory management comes with built-in DOS-UMB support. (page 4)

“AntiVirus detects and cleans both known and unknown viruses. (page 6)

“Lazy writing adds significant performance improvements to write operations. (page 8)

“To make compression really effective, it needs to be an integral part of the DOS file system. (page 15)

“When defragmentation is complete, all files are contiguous, and disk performance can increase noticeably. (page 18)

“The installation of the SuperStor/DS compression utilities is straightforward. (page 22)

“Workplace technologies are innovations that IBM will share across the OS/2, OS/400, and AIX products. (page 25)

“IBM Workplace will provide more natural user interfaces to the computer. (page 26)

“The IBM microkernel-based technology enables the operating system to be independent of the type of hardware processor. (page 30)